

THE COMBINATORICS OF MORSE THEORY WITH BOUNDARY

JONATHAN M. BLOOM

ABSTRACT. We prove several combinatorial results on path algebras over discrete structures related to directed graphs. These results are motivated by Morse theory on a manifold with boundary and, more generally, by Floer theory on a configuration space with boundary. Their purpose is to organize cobordism relationships among moduli spaces in order to define new algebraic invariants. We discuss applications to the Morse and Fukaya categories, and to work with John Baldwin on a bordered monopole Floer theory.

CONTENTS

1. Introduction	2
1.1. Counting paths	2
1.2. The view from Morse theory	4
1.3. Morse and Floer theory with boundary	6
2. Higraphs	10
2.1. Paths	10
2.2. Path operations	11
2.3. Path DGA	13
3. Abstract paths	15
3.1. Gluing moves on abstract paths	16
3.2. Abstract paths of low weight	18
4. Bigraphs	23
4.1. Paths	24
4.2. Path DGA	25
4.3. Tree-like higraphs and bigraphs	26
4.4. Balanced higraphs and bigraphs	27
4.5. Boundary bigraphs	28
4.6. Gluing sequences	28
Appendix: Morse homology with boundary	37
A. Morse homology on a closed manifold	37
B. Morse homology on a manifold with boundary	37
C. Relation to cell structures and CW-homology	40
References	43

The author was supported by NSF grant DMS-1104032.

1. INTRODUCTION

In this article, we prove several combinatorial results on path algebras over discrete structures related to directed graphs. These results are motivated by Morse theory on a manifold with boundary and, more generally, by Floer theory on a configuration space with boundary. Their purpose is to organize cobordism relationships among moduli spaces in order to define new algebraic invariants. Sections 2 through 4 give a fully self-contained exposition of the mathematical content. The Appendix briskly reviews Morse homology on a manifold with boundary, as defined by Kronheimer and Mrowka as the model for monopole Floer homology [9]. We also explain there the connection to cell structures and CW-homology. In the following three subsections of the introduction, we indicate the flavor of the combinatorics, its relation to Morse theory, and applications to Morse and Floer theory with boundary.

1.1. Counting paths. Let \mathbb{F} be the 2-element field. To a transitive digraph G , one can associate a differential graded algebra \mathcal{A} over \mathbb{F} , the *path DGA* of G . The \mathbb{F} -basis of \mathcal{A} is given by paths in G , the product by concatenation, and the grading by path length. The differential δ of the edge (v_1, v_2) is the sum of all length-two paths from v_1 to v_2 , and δ extends to paths by the Leibniz rule. The identity $\delta^2 = 0$ follows from the fact that a sequence of three compatible edges can be concatenated in two orders: $(ab)c$ and $a(bc)$. The element $D \in \mathcal{A}$ given by the sum of all the edges of G satisfies the *structure equation*

$$\delta D = D \circ D$$

as both sides equal the sum of all length-two paths in G . The digraph corresponding to Morse homology on a closed manifold is shown in Figure 1.

In a directed hypergraph, the source and target of an edge are each subsets of the set of vertices. In Section 2, we introduce a class of directed hypergraphs called *higraphs* on which the path DGA is well-defined and the structure equation continues to hold. The higraph underlying the cup product is shown in Figure 2.

In Sections 3 and 4 we extend these results to our main character: the *bigraph* of a higraph (the “b” stands for blow-up and boundary). Each higraph vertex is covered by three bigraph vertices called interior (o), stable (s), and unstable (u), and each higraph edge is covered by a number of bigraph edges called interior and boundary. The grading on the path DGA $\tilde{\mathcal{A}}$ of a bigraph is defined in terms of weight rather than length. The weight of an interior edge is 1 and the weight of a boundary edge depends on the number and types of vertices in its source and target (Definition 3.7). The differential $\tilde{\delta}$ on a weight w interior edge (resp., boundary edge) is the sum all weight $w + 1$ paths (resp., boundary paths) with the same source and target. Our first main results (Theorem 4.8) verifies that $\tilde{\delta}^2 = 0$. A bigraph is *tree-like* if every edge has singleton target and *rake-like* if every edge has singleton source. For tree-like bigraphs, our second result (Theorem 4.11) states that the element $\tilde{D} \in \tilde{\mathcal{A}}$ given by the sum of all weight 1 paths between interior and unstable vertices satisfies the structure equation

$$\tilde{\delta} \tilde{D} = \tilde{D} \circ \tilde{D}.$$

Dually, for rake-like bigraphs, the structure equation holds when \tilde{D} is given by the sum of all weight 1 paths between interior and stable vertices.




FIGURE 1. The higraph (and digraph) consisting of a single vertex and a single edge corresponds to Morse homology on a closed manifold. The vertex o represents the Morse complex and the edge D represents the Morse differential. The differential δ on the path DGA encodes the key cobordism relation between trajectory spaces illustrated on the surface at left. The structure equation is immediate. At right, the identity $\delta^2 = 0$ encodes the fact that each moduli space of twice-broken trajectories bounds two moduli spaces of once-broken trajectories.




FIGURE 2. The higraph at left consists of three vertices and four edges. The three edges labeled by μ_1 are loops and the fourth edge μ_2 runs from the two top vertices to the bottom vertex. In the context of the Morse category discussed in Section 1.3, the structure equation implies that μ_1 is a differential and μ_2 is a chain map, inducing the cup product structure.

The bigraph covering the higraph in Figure 1 is shown at the top of Figure 3. The reader may look to the Appendix, particularly Figure 15, to see how this fundamental example encodes the structure of Morse homology on a manifold with boundary. There are three vertices, four interior edges in black, and four boundary edges in red. The dashed red edge from s to u has weight 0 and the doubled red edge from u to s has weight 2. The other six edges have weight 1. The element D is the sum of all weight 1 paths starting or ending at o or s . The middle portion of Figure 3 shows the value of the differential $\tilde{\delta}$ on each edge, together with a surface

illustrating the cobordism relation between trajectory spaces encoded by $\tilde{\delta}(o, o)$. The bottom portion depicts the cancellation of three pairs of paths in $\tilde{\delta}\tilde{D}$ which do not appear in $\tilde{D} \circ \tilde{D}$. The greyscale surfaces along the bottom represent boundary components of 3-manifolds, in order to indicate the relationship of this cancellation to Morse theory. The families of red trajectories lie on the boundary while the families of black trajectories flow into the blue interior. The cancellation reflects the fact that each highlighted broken trajectory arises as the end of a one-parameter family of (broken) trajectories in two ways.

The bigraph in Figure 10 covers the higraph in Figure 2. Now $\tilde{\delta}\tilde{D}$ expands out to 197 paths, 62 of which cancel in pairs, leaving the 135 terms of $\tilde{D} \circ \tilde{D}$. A careful analysis of this cancellation phenomenon underlies the proof of Theorem 4.11.

1.2. The view from Morse theory. The Morse homology package extracts algebraic invariants of smooth manifolds from cobordism relationships between moduli spaces of gradient trajectories. The fundamental cobordism relation takes the schematic form

$$(1) \quad \partial X = X \circ X,$$

in which the boundary of one moduli space is identified with products of others moduli spaces. More precisely, we have

$$(2) \quad \partial \bar{X}_b^a = \bigcup_c X_c^a \times X_b^c.$$

where X_b^a is the moduli space of unparameterized trajectories from a to b , and \bar{X}_b^a denotes a compactification of X_b^a . One defines an endomorphism D on the vector space generated by the set of critical points by setting

$$\langle Da, b \rangle = |X_b^a|.$$

One may also define an endomorphism δD by setting

$$\langle (\delta D)a, b \rangle = |\partial \bar{X}_b^a|.$$

The cobordism relation (2) implies the *structure equation*

$$(3) \quad \delta D = D \circ D,$$

a precise algebraic translation of the schematic (1). Furthermore, δD is the zero map, since $\partial \bar{X}_b^a$ is finite if and only if it is the boundary of a compact 1-manifold and hence an even number of points. We conclude from the structure equation that D is indeed a differential. Note that the same framework is used to prove that Morse homology is invariant and functorial. Going one step further, one may define an endomorphism $\delta\delta D$ by setting

$$\langle (\delta\delta D)a, b \rangle = |\partial(\partial \bar{X}_b^a)|.$$

This is the zero map for a different topological reason: the compactification of X_b^a may be given the structure of a smooth manifold with corners.

This viewpoint naturally extends to the case of a manifold with boundary.

For a manifold with boundary, the space X is itself a union of products of moduli spaces of trajectories (that is, X is composed of several types of possibly-broken trajectories). The space \bar{X} is built from the compactifications of the factors, its boundary given by the Leibniz

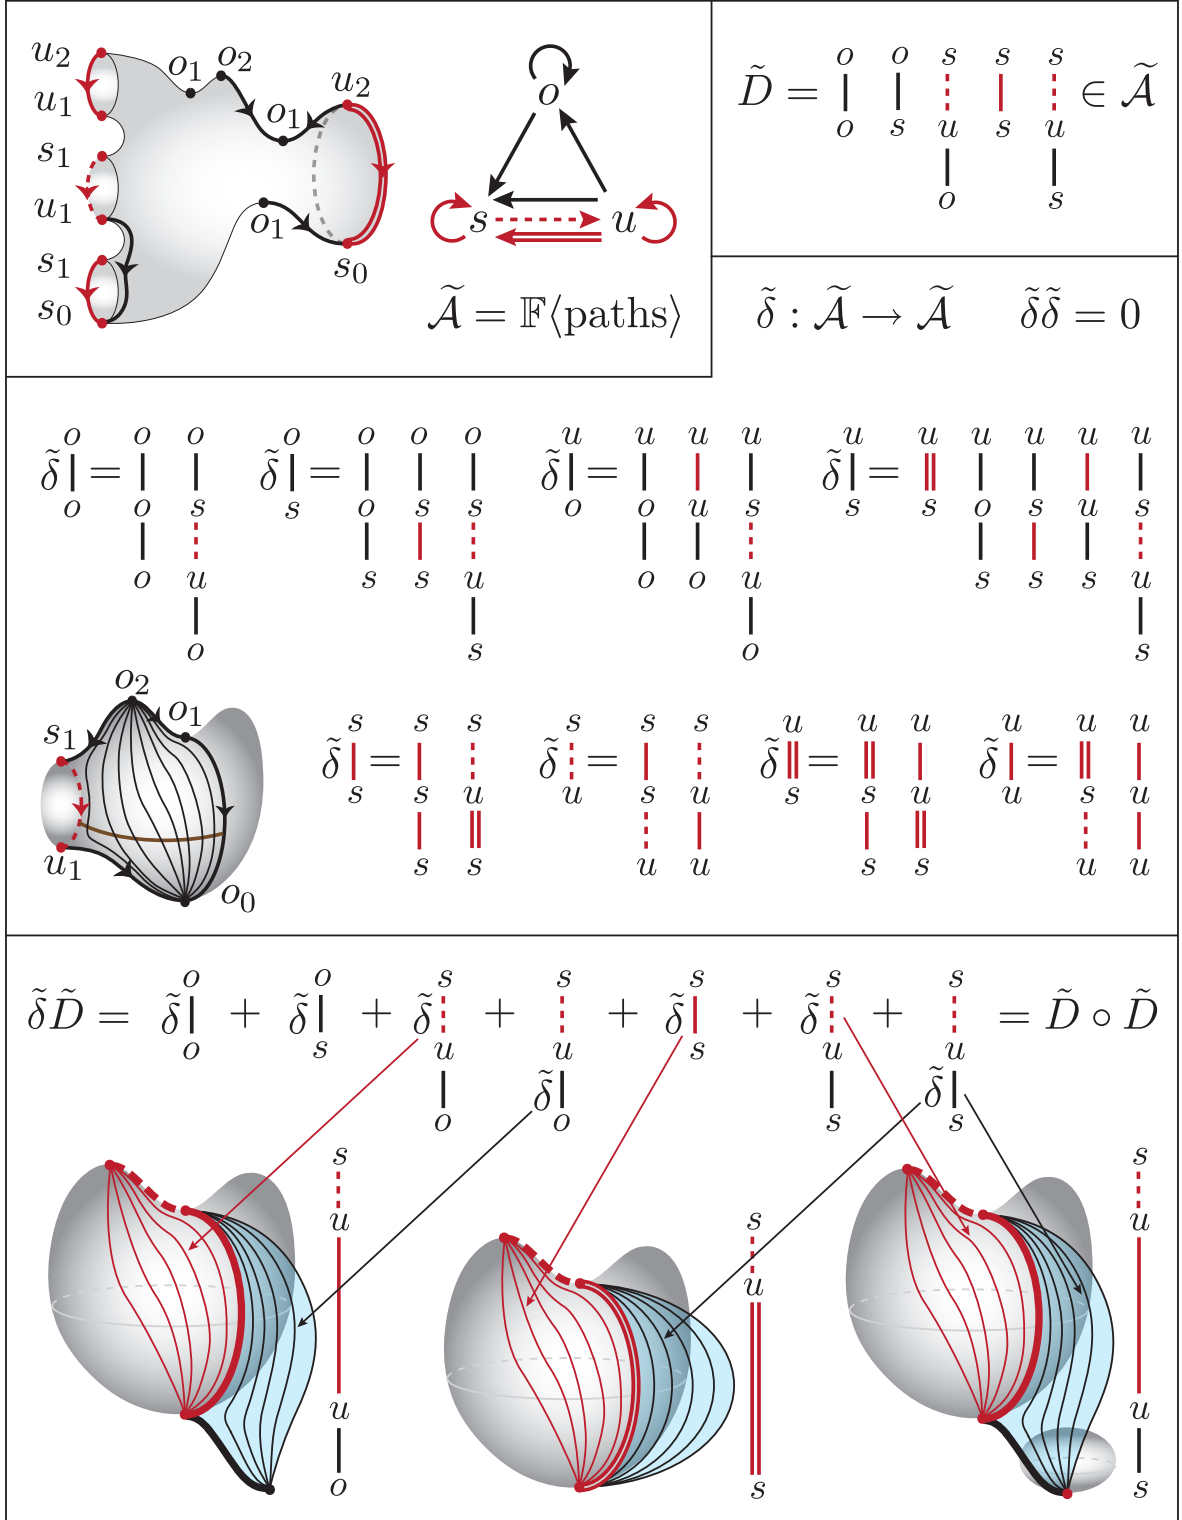


FIGURE 3. The path DGA of the bigraph representing Morse homology on a manifold with boundary (compare with Figure 15).

rule. The structure equation continues to hold, although now after the cancellation of several pairs of terms in δD which do not appear in $D \circ D$ as illustrated in Figure 3. And the maps δD and $\delta\delta D$ still vanish. We can summarize the situation as follows: without boundary, Morse homology depends on geometric analysis and a bit of topology; with boundary, Morse homology depends on (more involved) geometric analysis, a bit of topology, and the simplest case of the combinatorics that is comprehensively developed herein.

Remark 1.1. Kronheimer and Mrowka study Morse theory with boundary in the context of monopole Floer homology. They do not prove that \overline{X} is a manifold with corners, but rather that it is stratified by smooth manifolds and has an even number of ends in the 1-dimensional case. The example in Figure 12 indicates that \overline{X} is not a smooth manifold with corners when its dimension is three or more. On the other hand, Proposition 4.22 shows that $|\partial(\partial\overline{X})| = 0$, which corresponds to the identity $\delta^2 = 0$ in the path DGA setting and suggests that \overline{X} may be a manifold with corners in the 2-dimensional case.

1.3. Morse and Floer theory with boundary. Morse homology is an approach to analyzing the topology of a smooth manifold by studying the gradient flow of a generic smooth functional. Floer theory generalizes this approach to infinite-dimensional configuration spaces arising in symplectic geometry and gauge theory. Two important examples, Lagrangian intersection Floer homology and monopole Floer homology, may be placed in parallel with finite-dimensional Morse homology as follows:

- The *Morse complex* $C(f)$ of a smooth function f on a closed manifold is generated by the critical points of f ; the differential counts gradient trajectories between critical points.
- The *Langrangian Floer complex* $CF(L_0, L_1)$ of two Lagrangian submanifolds of a symplectic manifold X is generated by the points of $L_0 \cap L_1$; the differential counts pseudo-holomorphic bigons with edge i on L_i . Here the functional is the symplectic action functional on the space of paths from L_0 to L_1 [8].
- The *monopole Floer complex* $\widehat{C}(Y)$ of a closed 3-manifold Y is generated by solutions (called monopoles) to the 3-dimensional Seiberg-Witten equations on Y ; the differential counts solutions to the 4-dimensional Seiberg-Witten equations on $Y \times \mathbb{R}$. Here the functional is the Chern-Simons-Dirac functional on the configuration space $\mathcal{B}(Y)$ consisting of pairs of certain connections and spinors over Y modulo gauge [9].

We caution that these rough descriptions suppresses auxiliary data (such as metrics, perturbations, and almost-complex structures) and ignore bubbling phenomena that may prevent the Langrangian Floer complex from being well-defined. Note also that the configuration space $\mathcal{B}(Y)$ has a boundary consisting of reducible configurations, which is why there are three versions of monopole Floer homology. The correspondence with Morse homology is reflected by the exact sequences

$$\begin{aligned} \cdots \longrightarrow H_*(\partial M) \longrightarrow H_*(M) \longrightarrow H_*(M, \partial M) \longrightarrow \cdots \\ \cdots \longrightarrow \overline{HM}_\bullet(Y) \longrightarrow \widetilde{HM}_\bullet(Y) \longrightarrow \widehat{HM}_\bullet(Y) \longrightarrow \cdots \end{aligned}$$

Singular cohomology admits an associative product structure, the cup product, which may be defined Morse-theoretically using gradient trees (or equivalently, triple intersections of

stable and unstable manifolds with respect to several Morse functions). Fukaya recognized that deeper multiplicative structure is naturally encoded by an A_∞ -category, which determines the Massey products and rational homotopy type [7].

- The objects of the *Morse category* of a closed manifold are smooth functions and $\text{Mor}(f_0, f_1) = C(f_0 - f_1)$. The multiplication map

$$\mu_k : C(f_0 - f_1) \otimes \cdots \otimes C(f_{k-1} - f_k) \rightarrow C(f_0 - f_k)$$

counts gradient trees with k inputs and 1 output.

Here μ_1 is the Morse differential and μ_2 is the chain map defined by

$$\langle \mu_2(a \otimes b), c \rangle = |\{x \in M : \lim_{t \rightarrow \infty} \phi_t^{f_0 - f_1}(x) = a, \lim_{t \rightarrow \infty} \phi_t^{f_1 - f_2}(x) = b, \lim_{t \rightarrow \infty} \phi_t^{f_0 - f_2}(x) = c\}|$$

where $\phi_t^f : M \rightarrow M$ denotes the flow induced by ∇f . The map μ_3 provides a homotopy between $\mu_2(\mu_2(\cdot, \cdot), \cdot)$ and $\mu_2(\cdot, \mu_2(\cdot, \cdot))$. In particular, μ_2 induces an associative product on homology, dual to the cup product. See [15] for background on A_∞ -categories.

One application of our combinatorial results is to extend the definition of the Morse category to manifolds with boundary. For the appropriate sequence of higraphs, the path DGAs of bigraphs specify the higher multiplications, and the structure equation is equivalent to the A_∞ -relations. For instance, the map μ_2 is defined to count broken gradient trees as specified in Figure 10. The higraph for μ_3 is shown in Figure 5, and the general case is similar.

Fukaya introduced the Morse category as a model for his eponymous category, which encodes the higher multiplicative structure of Langrangian intersection Floer homology.

- The objects of the *Fukaya category* of a symplectic manifold X are Lagrangian submanifolds and $\text{Mor}(L_0, L_1) = CF(L_0, L_1)$. The multiplication map

$$\mu_k^{\text{Lag}} : CF(L_0, L_1) \otimes \cdots \otimes CF(L_{k-1}, L_k) \rightarrow CF(L_0, L_k)$$

counts pseudo-holomorphic $(k+1)$ -gons with edge i on L_i .

The connection with the Morse category goes as follows. Given a smooth function f on M , the graph of df is a Lagrangian in T^*M with the standard symplectic structure. The Fukaya category generated by such Lagrangians is A_∞ -equivalent to the Morse category of M , which implies the Arnold conjecture for the zero-section [7], [1].

The Fukaya category is a central algebraic structure in symplectic geometry and mirror symmetry. It is also notoriously difficult to define or compute rigorously [8], [15]. See [13] for a state-of-the-art computation: the Fukaya category of the 2-dimensional torus. A divide-and-conquer approach to computation might go as follows: chop the symplectic manifold into simpler pieces and reassemble its Fukaya category by algebraically compiling the data of some notion of Fukaya category for each piece. Akaho has taken a step in this direction, defining the Floer complex $CF(L_1, L_2)$ of two Lagrangians in an open symplectic manifold with concave end [2]. A slice of the end inherits a contact structure and intersects the Lagrangians in Legendrians. The complex $CF(L_1, L_2)$ is generated by both points of $L_1 \cap L_2$ and Reeb chords between the Legendrians. The Reeb chords play the role of boundary critical points and the differential takes the same form as the Morse differential on a manifold with boundary. As such, our bigraph framework should specify the higher multiplications needed to define the Fukaya category of an open symplectic manifold with concave end as well.

Our primary application, in joint work with John Baldwin, is to define a gauge-theoretic analogue of the Fukaya category using monopole Floer theory. This monopole category combines older ideas of Segal, Donaldson, and Fukaya [6] with the comprehensive analysis of Kronheimer and Mrowka [9].

- The objects of the *monopole category* of a closed surface Σ are 3-manifolds with boundary parameterized by Σ , and $\text{Mor}(Y_0, Y_1) = \widehat{C}(Y_0 \cup_{\Sigma} -Y_1)$. The map

$$\mu_k^{\text{mon}} : \widehat{C}(Y_0 \cup -Y_1) \otimes \cdots \otimes \widehat{C}(Y_{k-1} \cup -Y_k) \rightarrow \widehat{C}(Y_0 \cup -Y_k)$$

counts monopoles on a 4-dimensional cobordism

$$W : (Y_0 \cup -Y_1) \sqcup \cdots \sqcup (Y_{k-1} \cup -Y_k) \rightarrow Y_0 \cup -Y_k$$

over a family of metrics parameterized by the $(k-2)$ -dimensional associahedron.

The cobordism W is constructed by gluing $\partial Y_i \times [0, 1] \subset Y_i \times [0, 1]$ to $\Sigma \times e_i \subset \Sigma \times \Delta$ for $0 \leq i \leq k$, where Δ is a polygonal region bounded by $k+1$ line segments alternating with $k+1$ arcs e_i . The case $k=2$ is depicted at right in Figure 4. In general, the family of metrics interpolates between all ways of cutting W along disjoint hypersurfaces $Y_i \cup -Y_j$ corresponding to chords in Δ between non-adjacent arcs e_i . Figure 5 illustrates the interval of metrics in the case $k=3$ from a less rigid perspective.

In light of the isomorphism between monopole and Heegaard Floer homology (see [4] and [11]), it is natural to conjecture that the monopole category of a surface Σ of genus g is A_{∞} -equivalent to the subcategory of the Fukaya category of $\text{Sym}^g(\Sigma)$ generated by Heegaard tori. Furthermore, the monopole category fits into the framework of a *bordered monopole Floer theory*, a form of extended TQFT. Roughly, the bordered monopole Floer package is a 2-functor sending surfaces to A_{∞} -categories, 3-dimensional cobordisms to A_{∞} -functors, and 4-dimensional cobordisms with corners to A_{∞} -natural transformations. In dimensions 2 and 3, this theory is analogous to the version of bordered Heegaard Floer theory developed by Lekili and Perutz [12] (which in turn is related to the original version [14]). In our case, for each genus, the monopole category is split-generated by an explicit, finite set of parameterized handlebodies through the surgery exact triangle, and the homology algebra is known. We hope




FIGURE 4. The μ_2 multiplication map in the Morse, Fukaya, and monopole A_{∞} -categories counts gradient trees, pseudo-holomorphic triangles, and monopoles on a 4-dimensional cobordism, respectively. On homology, this map induces composition of morphisms in an ordinary category.




FIGURE 5. The μ_3 multiplication map is represented by the green edge from the three top vertices to the bottom vertex. The full higraph contains 6 vertices and 11 edges. The compositions $\mu_2(\mu_2(\cdot, \cdot), \cdot)$ and $\mu_2(\cdot, \mu_2(\cdot, \cdot))$ are each represented by paths of length 2 at right. Due to boundary (reducibles), the terms of the monopole maps μ_1^{mon} , μ_2^{mon} , μ_3^{mon} are encoded by the corresponding bigraph. Roughly, μ_3^{mon} counts monopoles on a 4-dimensional cobordism (at left) over a family of metrics parameterized by the 1-dimensional associahedron: the interval $[-1, 1]$. The metric stretches normal to the red (resp., blue) hypersurface as the parameter goes to -1 (resp., 1). The degenerate metrics at ± 1 correspond to composite cobordisms (at right). The structure equation implies the A_∞ -relation for μ_3^{mon} (not shown).

that this theory will point toward an axiomatization of Floer homology in low-dimensional topology.

The preceding applications utilize the sequence of higraphs arising from associahedra encoding homotopy associativity. Other higraphs are useful as well. The link surgery spectral sequence is most naturally defined using permutohedra encoding homotopy commutativity; these higraphs have the vertex set of a hypercube with an edge between every ordered pair of vertices [3]. Baldwin and the author's constructions in bordered monopole Floer theory make use of more general polytopes called graph associahedra [5], including cyclohedra related to Hochschild homology. More basically, a tree higraph (one non-loop edge with singleton target) leads to an \widehat{HM}_\bullet map for a connected cobordism with disconnected source as in Figure 10. Dually, a rake higraph (one non-loop edge with singleton source) leads to an \widehat{HM}_\bullet map for a connected cobordism with disconnected target. As a final example, we have appended the higraph organizing the proof of the monopole surgery exact triangle as Figure 19. To cover these and future cases, we have aimed for a definition of higraphs that is as general as possible while still satisfying a unified combinatorial theory.

2. HIGRAPHS

Definition 2.1. A *directed graph*, or *digraph*, $G = (V, E)$ consists of a vertex set V and an edge set $E \subset V \times V$. The edge $e = (v_1, v_2)$ has *source* $s(e) = v_1$ and *target* $t(e) = v_2$.

A *directed hypergraph* $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ consists of a vertex set \mathcal{V} and an edge set $\mathcal{E} \subset 2^{\mathcal{V}} \times 2^{\mathcal{V}}$. The edge $\varepsilon = (I, J)$ has *source* $s(\varepsilon) = I \subset \mathcal{V}$ and *target* $t(\varepsilon) = J \subset \mathcal{V}$.

We reserve the term higraph for a special kind of directed hypergraph.

Definition 2.2. A *higraph* \mathcal{H} is a directed hypergraph $(\mathcal{V}, \mathcal{E})$ such that the following properties hold for all edges $\varepsilon, \varepsilon' \in \mathcal{E}$.

TRANSITIVE: If $t(\varepsilon) \cap s(\varepsilon')$ is non-empty, then $t(\varepsilon) \cap s(\varepsilon') = \{K\}$ for some $K \in \mathcal{V}$ and

$$(4) \quad (s(\varepsilon) \cup (s(\varepsilon') - \{K\}), (t(\varepsilon) - \{K\}) \cup t(\varepsilon')) \in \mathcal{E}$$

where each of the above unions is of disjoint sets.

ACYCLIC: If $t(\varepsilon) \cap s(\varepsilon')$ and $t(\varepsilon') \cap s(\varepsilon)$ are both non-empty, then $\varepsilon = \varepsilon' = (\{K\}, \{K\})$ for some $K \in \mathcal{V}$.

An edge of the form $(\{K\}, \{K\})$ is called a *loop*. Transitivity ensures that adjacent edges ε and ε' may be glued to form a single edge (4) in a well-defined manner. Acyclicity ensures that the only edges with overlapping source and target are loops, and that loops cannot result from gluing distinct adjacent edges.




FIGURE 6. Of the nine directed hypergraphs above, only the four at right are higraphs. All vertices are shown as black dots. The bottom-right higraph has only one edge; the source and target of this edge each consist of three vertices. Two edges in the top-right higraph have empty source; gluing the upper two edges results in the lower edge.

2.1. Paths. We recall the notion of a directed tree.

Definition 2.3. A *directed tree* T is a digraph which is acyclic as an undirected graph. For clarity, we will refer to the vertices and edges of T as *nodes* and *arrows*, respectively.

A path in a higraph has the structure of a directed tree – the nodes index the edges of the path, while the arrows index the breaks between edges.

Definition 2.4. A *path* $\rho = (T, \rho)$ in a higraph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ consists of a non-empty, directed tree $T = (V, E)$ together with a map $\rho : V \rightarrow \mathcal{E}$ such that:

(1) The map ρ induces a well-defined map $\rho : E \rightarrow \mathcal{V}$ such that

$$\{\rho(e)\} = t(\rho(s(e))) \cap s(\rho(t(e))).$$

(2) If distinct arrows e and e' have the same source or the same target, then $\rho(e) \neq \rho(e')$.

Paths ρ and ρ' are *isomorphic* if there is a digraph isomorphism $f : T \rightarrow T'$ of the associated directed trees such that $\rho = \rho'f$.

In fact, the automorphism group of a path is trivial. This is clear if ρ covers a loop. If ρ does not cover a loop, then transitivity and acyclicity imply that the leaves of T are mapped to distinct edges. So in this case, any automorphism of ρ must fix the leaves of T , and is therefore the identity. Thus, paths are unique up to unique isomorphism, and we consider isomorphic paths to be the same.

Definition 2.5. Let \mathcal{P} denote the set of paths in \mathcal{H} . Given a path ρ , the *length* of ρ is the number of edges it contains, given by $l(\rho) = |V|$. We refer to the arrows of T as *breaks*. The number of breaks in ρ is given by $b(\rho) = |E| = l(\rho) - 1$. The *source* of a path ρ is the subset of \mathcal{V} given by

$$s(\rho) = \bigcup_{v \in V} (s(\rho(v)) - \{\rho(e) \mid t(e) = v\}).$$

The *target* of a path ρ is the subset of \mathcal{V} given by

$$t(\rho) = \bigcup_{v \in V} (t(\rho(v)) - \{\rho(e) \mid s(e) = v\}).$$

Transitivity implies that the union defining $s(\rho)$ is of disjoint sets, and similarly for $t(\rho)$. Thus each vertex $K \in s(\rho)$ is *sourced* by a unique node $v \in V$. Similarly, each vertex $K \in t(\rho)$ is *targeted* by a unique node $v \in V$.

2.2. Path operations. We next define several ways to combine or modify paths. *Concatenation* joins two compatible paths together. *Swapping* cuts out a subpath of a path and replaces it with another compatible path. *Refinement* and *gluing* are special cases of swapping. Refinement expands a single edge to a compatible path, while gluing contracts a subpath to a single edge.

2.2.1. Concatenation.

Definition 2.6. The partially-defined *concatenation* map $c : \mathcal{P} \times \mathcal{P} \dashrightarrow \mathcal{P}$ joins compatible paths together. Consider paths ρ_1 and ρ_2 with associated trees $T_1 = (V_1, E_1)$ and $T_2 = (V_2, E_2)$. The path $\rho = c(\rho_1, \rho_2)$ is defined precisely when $t(\rho_1) \cap s(\rho_2)$ is non-empty, in which case the intersection contains a single vertex $K \in \mathcal{V}$ targeted by $v_1 \in V_1$ and sourced by $v_2 \in V_2$. The directed tree $T = (V, E)$ associated to ρ is given by

$$V = V_1 \cup V_2$$

and

$$E = E_1 \cup E_2 \cup \{(v_1, v_2)\}.$$

The map $\rho : V \rightarrow \mathcal{E}$ is given by $\rho_1 \sqcup \rho_2$. The induced map $\rho : E \rightarrow \mathcal{V}$ sends (v_1, v_2) to K .

2.2.2. *Swapping.* Let $\mathcal{Q} = \{(\rho, T') \mid \rho = (T, \rho) \in \mathcal{P}, T' \text{ is a non-empty subtree of } T\}$.

Definition 2.7. The partially-defined *swapping map* $w : \mathcal{Q} \times \mathcal{P} \dashrightarrow \mathcal{P}$ replaces a subpath of a path with a compatible path. Given a path $\rho_1 = (T_1, \rho_1)$ and non-empty subtree $T' = (V', E')$ of T_1 , we obtain a subpath $\rho' = (T', \rho')$ by setting $\rho' = \rho_1|_{V'}$. The path $\rho_2 = w((\rho_1, T'), \rho'')$ is defined precisely when $s(\rho') = s(\rho'')$ and $t(\rho') = t(\rho'')$. In this case, the associated tree $T_2 = (V_2, E_2)$ has node set given by

$$V_2 = (V_1 - V') \cup V''.$$

To describe the arrow set E_2 , note that for each arrow $e \in (V_1 - V') \times V'$, the vertex $\rho'(t(e)) \in s(\rho'')$ is sourced by a unique node $t''(e)$ of T'' . Similarly, for each arrow $e \in V' \times (V_1 - V')$, the vertex $\rho'(s(e)) \in t(\rho'')$ is targeted by a unique node $s''(e)$ of T'' . We set

$$E_2 = (E_1 \cap ((V_1 - V') \times (V_1 - V'))) \cup E_s \cup E_t \cup E'' \subset V_2 \times V_2$$

where

$$E_s = \{(s(e), t''(e)) \mid e \in (V_1 - V') \times V'\}$$

and

$$E_t = \{(s''(e), t(e)) \mid e \in V' \times (V_1 - V')\}.$$

The map $\rho : V_2 \rightarrow \mathcal{E}$ is induced by ρ_1 on $V_1 - V'$ and by ρ'' on V'' . Note that ρ_2 has the same source and target as the original path ρ_1 .

Since swapping preserves source and target, swapping and concatenation commute in the following sense. Fix paths ρ_1 , ρ_2 , and ρ'' , and let T'_1 and T'_2 be non-empty subtrees of T_1 and T_2 , respectively. Then

$$(5) \quad w((c(\rho_1, \rho_2), T'_1), \rho'') = c(w((\rho_1, T'_1), \rho''), \rho_2)$$

$$(6) \quad w((c(\rho_1, \rho_2), T'_2), \rho'') = c(\rho_1, w((\rho_2, T'_2), \rho''))$$

where the left-hand side of each equation is defined if and only if the right-hand side is defined.

Of course, if we swap out a subpath and then swap it back in, we end up where we started. More precisely, if we define the map $W : \mathcal{Q} \times \mathcal{P} \rightarrow \mathcal{Q} \times \mathcal{P}$ by

$$(7) \quad W((\rho_1, T'), \rho'') = ((w((\rho_1, T'), \rho''), T''), \rho'),$$

then $W \circ W$ is the identity map. We are most interested in the following two special cases of swapping, which are mutual inverses in the sense of (7).

2.2.3. *Refinement.* We first suppose that the subpath ρ' consists of a single edge $\varepsilon^* = \rho'(v^*)$. Let

$$\mathcal{R} = \{(\rho, v^*) \mid \rho = (T, \rho) \in \mathcal{P}, v^* \text{ is a vertex of } T\}.$$

Definition 2.8. The partially-defined *refinement map* $r : \mathcal{R} \times \mathcal{P} \dashrightarrow \mathcal{P}$ replaces a single edge $\varepsilon^* = \rho_1(v^*)$ of a path ρ_1 with a compatible path ρ'' :

$$r((\rho_1, v^*), \rho'') = w((\rho_1, v^*), \rho'').$$

2.2.4. *Gluing.* We next suppose that the substituted path ρ'' consists of a single edge ε^* .

Definition 2.9. The *gluing map* $q : \mathcal{Q} \rightarrow \mathcal{P}$ replaces a subpath ρ' of a path ρ_1 with the single edge $\varepsilon^* = (s(\rho'), t(\rho'))$:

$$q(\rho_1, T') = w((\rho_1, T'), \varepsilon^*).$$

Note that ε^* is indeed an edge of \mathcal{H} by transitivity.

By definition, the tree T_2 underlying $\rho_2 = q(\rho_1, T')$ results from contracting the subtree T' of T_1 to a single vertex v^* , with $\rho_2(v^*) = \varepsilon^*$. In particular, the projection map $\pi : V_1 \rightarrow V_2$ given by

$$(8) \quad \pi(v) = \begin{cases} v & \text{if } v \in V_1 - V', \\ v^* & \text{if } v \in V'. \end{cases}$$

induces a bijection $\pi : (E_1 - E') \rightarrow E_2$ between the arrows of T_1 that are outside T' and the arrows of T_2 .

Definition 2.10. Let e be a break in a path ρ . We use the notation $q(\rho, e)$ to denote the result of gluing the path ρ along the one-arrow subtree induced by e . Note that gluing a path at a break reduces length by 1.

Definition 2.11. A *k-step gluing sequence* on a path ρ_1 is a k -tuple (e_1, \dots, e_k) with e_i a break on ρ_i with $\rho_{i+1} = q(\rho_i, e_i)$ for $1 \leq i \leq k$. We may write this as

$$\rho_1 \xrightarrow{e_1} \rho_2 \xrightarrow{e_2} \dots \xrightarrow{e_k} \rho_{k+1}$$

to emphasize the paths along the way. An isomorphism of k -step gluing sequences is a sequence of isomorphisms between the corresponding paths which identify the corresponding breaks.

Gluing sequences are unique up to unique isomorphism, so we consider isomorphic gluing sequences to be the same. The following lemma is established in Section 4.6.1.

Lemma 2.12. *The number of k-step gluing sequences between any two fixed paths in a higraph is a multiple of $k!$.*

2.3. **Path DGA.** The path algebra on a digraph is associative but not commutative. The following generalization to higraphs is neither associative nor commutative in general.

Definition 2.13. The *path DGA* of the higraph \mathcal{H} over the field \mathbb{F} consists of a graded algebra \mathcal{A} equipped with a differential $\delta : \mathcal{A} \rightarrow \mathcal{A}$. The algebra \mathcal{A} has underlying vector space $\mathbb{F}\langle \mathcal{P} \rangle$, with the subspace \mathcal{A}^k in grading k spanned by the set \mathcal{P}^k of paths of length k . Multiplication is bilinear and given on paths by $\rho_1 \rho_2 = c(\rho_1, \rho_2)$ if defined, and $\rho_1 \rho_2 = 0$ otherwise.

The differential δ is determined by the Leibniz rule and linearity once we specify its value on each edge $\varepsilon \in \mathcal{E}$ regarded as an element of \mathcal{A} . The value of δ on an edge ε equals the sum of all refinements of ε to a path of length 2. In particular, the differential has degree 1.

Since a path ρ may often be expressed as a fully-associated product of edges in several ways, we should take care that δ is well-defined. This issue is resolved by noting that we may define δ directly on a path $\rho = ((V, E), \rho)$ by

$$(9) \quad \delta \rho = \sum_{v \in V} \sum_{\rho' \in \mathcal{P}^2} r((\rho, v), \rho')$$

where $r((\rho, v), \rho')$ is defined to be zero on invalid input. This definition is consistent with our earlier definition of $\delta\varepsilon$, and the Leibniz rule now follows from (5) and (6):

$$\begin{aligned}
\delta(\rho_1\rho_2) &= \sum_{v \in V} \sum_{\rho' \in \mathcal{P}^2} r((\rho_1\rho_2, v), \rho') \\
&= \sum_{v \in V_1} \sum_{\rho' \in \mathcal{P}^2} r((\rho_1\rho_2, v), \rho') + \sum_{v \in V_2} \sum_{\rho' \in \mathcal{P}^2} r((\rho_1\rho_2, v), \rho') \\
&= \left(\sum_{v \in V_1} \sum_{\rho' \in \mathcal{P}^2} r((\rho_1, v), \rho') \right) \rho_2 + \rho_1 \left(\sum_{v \in V_2} \sum_{\rho' \in \mathcal{P}^2} r((\rho_2, v), \rho') \right) \\
&= (\delta\rho_1)\rho_2 + \rho_1(\delta\rho_2).
\end{aligned}$$

To see that δ is a differential, it is useful to introduce its adjoint ∂ with respect to the symmetric, bilinear pairing

$$\langle \cdot, \cdot \rangle : \mathcal{A} \times \mathcal{A} \rightarrow \mathbb{F}$$

with $\langle \rho_1, \rho_2 \rangle = 1$ if and only if $\rho_1 = \rho_2$.

Definition 2.14. The map $\partial : \mathcal{A} \rightarrow \mathcal{A}$ is defined on the path ρ with underlying tree $T = (V, E)$ by

$$\partial\rho = \sum_{e \in E} q(\rho, e).$$

Proposition 2.15. *We have*

$$\langle \partial a, b \rangle = \langle a, \delta b \rangle$$

for all $a, b \in \mathcal{A}$. Furthermore, $\partial^2 = 0$ and $\delta^2 = 0$.

Proof. For the first statement, by linearity it suffices to verify that

$$\langle \partial\rho_1, \rho_2 \rangle = \langle \rho_1, \delta\rho_2 \rangle.$$

The left-hand side counts the number of 1-step gluing sequences $\rho_1 \xrightarrow{e_1} \rho_2$. The right-hand side counts all ways to refine an edge of ρ_2 to a path of length 2 to obtain ρ_1 . Equality follows from the fact that gluing and refinement are inverses as in (7).

To prove $\partial^2 = 0$, it suffices to show that $\langle \partial^2\rho_1, \rho_3 \rangle$ vanishes for all paths ρ_1 and ρ_3 . Indeed, the term $\langle \partial^2\rho_1, \rho_3 \rangle$ counts the number of 2-step gluing sequences $\rho_1 \xrightarrow{e_1} \rho_2 \xrightarrow{e_2} \rho_3$, which is even by Lemma 2.12.

Now $\delta^2 = 0$ because $\langle \delta^2 a, b \rangle = \langle \delta a, \partial b \rangle = \langle a, \partial^2 b \rangle = \langle a, 0 \rangle = 0$. \square

Definition 2.16. We define the element $D \in \mathcal{A}$ by

$$D = \sum_{\varepsilon \in \mathcal{E}} \varepsilon = \sum_{\rho \in \mathcal{P}^1} \rho.$$

Proposition 2.17. *The element D satisfies the structure equation*

$$\delta D = D \circ D.$$

Proof. We claim

$$\delta D = \sum_{\rho \in \mathcal{P}^2} \rho = D \circ D.$$

The left-hand equality holds because each path ρ of length 2 arises as the refinement of a unique edge $\varepsilon = (s(\rho), t(\rho))$. The right-hand equality holds because each path of length 2 arises as the product of two (not necessarily distinct) edges in a unique way. \square

3. ABSTRACT PATHS

In Section 4, we will define the notion of the bigraph $\tilde{\mathcal{H}}$ associated to a higraph \mathcal{H} . We will then systematically extend the arguments and results of Section 2 to bigraphs. In this section, we assemble the necessary lemmas, which concern the abstract form of a path in a bigraph. We call such a form an abstract path, as it is defined without reference to a particular bigraph. We now make this notion precise.

Definition 3.1. An *abstract vertex* is an element of the set $\mathcal{V}_{ab} = \{o, u, s\}$. The symbols o , u , and s stand for *interior*, *unstable*, and *stable*, respectively. Both unstable and stable abstract vertices are considered *boundary*.

An *abstract edge* is an element

$$\varepsilon = (\mu, (o_1, u_1, s_1), (o_2, u_2, s_2))$$

of the subset

$$\mathcal{E}_{ab} \subset \mathbb{F} \times (\mathbb{N} \times \mathbb{N} \times \mathbb{N}) \times (\mathbb{N} \times \mathbb{N} \times \mathbb{N}).$$

This subset is cut out by the following two conditions:

Condition 0: If $\mu = 0$ then $s_1 = u_2 = 0$.

Condition 1: If $\mu = 1$ then $o_1 = o_2 = 0$.

An abstract edge ε is *interior* if $\mu = 0$ and *boundary* if $\mu = 1$.

We think of the abstract edge ε as having $o_1 + u_1 + s_1$ incoming ends (or source vertices), of which o_1 are interior, u_1 are unstable, and s_1 are stable. Similarly, ε has $o_2 + u_2 + s_2$ outgoing ends, of which o_2 are interior, u_2 are unstable, and s_2 are stable. Condition 0 says that the incoming ends of an interior abstract edge are interior or unstable, while the outgoing ends are interior or stable. Condition 1 says that the ends of a boundary abstract edge are boundary abstract vertices.

We can assemble abstract edges into an abstract path, so long as there are sufficiently many ends for concatenation.

Definition 3.2. An *abstract path* is a triple $\tau = (T, \tau, \sigma)$ consisting of a non-empty directed tree T and maps

$$\begin{aligned} \tau : V &\rightarrow \mathcal{E}_{ab} \\ \sigma : E &\rightarrow \mathcal{V}_{ab} \end{aligned}$$

such that the following inequalities hold for each $v \in V$, with $\tau(v) = (\mu, (o_1, u_1, s_1), (o_2, u_2, s_2))$:

$$\begin{aligned} o_1 &\geq |\{e \in E \mid t(e) = v, \sigma(e) = o\}|, \\ u_1 &\geq |\{e \in E \mid t(e) = v, \sigma(e) = u\}|, \\ s_1 &\geq |\{e \in E \mid t(e) = v, \sigma(e) = s\}|, \\ o_2 &\geq |\{e \in E \mid s(e) = v, \sigma(e) = o\}|, \\ u_2 &\geq |\{e \in E \mid s(e) = v, \sigma(e) = u\}|, \\ s_2 &\geq |\{e \in E \mid s(e) = v, \sigma(e) = s\}|. \end{aligned}$$

An abstract path τ is *boundary* if all of its abstract edges are boundary.

We interpret an arrow e from node v_1 to node v_2 to signify that the abstract edges $\tau(v_1)$ and $\tau(v_2)$ are concatenated at an abstract vertex of type $\sigma(e)$. In this sense, an abstract path has a well-defined number of source and target abstract vertices of each type. For example, the number of incoming stable ends is given by

$$s_1(\tau) = \left(\sum_{v \in V} s_1(v) \right) - |\sigma^{-1}(s)|,$$

while the number of outgoing unstable ends is given by

$$u_2(\tau) = \left(\sum_{v \in V} u_2(v) \right) - |\sigma^{-1}(u)|.$$

Definition 3.3. An abstract path τ is *legal* if $s_1(\tau) = u_2(\tau) = 0$; equivalently, τ is legal if its ends are compatible with an interior abstract edge (see Condition 0 of Definition 3.1).

It is convenient to pull back the terminology in Definition 3.1, so that an arrow $e \in E$ is interior, unstable, or stable according to the value of $\sigma(e)$. Similarly, a node $v \in V$ is interior or boundary according to the value of $\mu(v)$. We may decompose the node set V as $V^0 \cup V^1$ where $V^0 = \mu^{-1}(0)$ denotes the set of interior nodes and $V^1 = \mu^{-1}(1)$ denotes the set of boundary nodes.

Definition 3.4. The *interior neighborhood* of a node $v \in V$ is the set of adjacent interior nodes:

$$\text{Nbhd}^o(v) = V^0 \cap \{v' \in V \mid (v, v') \in E \text{ or } (v', v) \in E\}.$$

The *interior star* of v is the subtree $\text{Star}^o(v)$ of T induced by $\{v\} \cup \text{Nbhd}^o(v)$.

3.1. Gluing moves on abstract paths. We have introduced all the necessary terminology to define the gluing moves on abstract paths. We illustrate these moves in Figure 7, while precise definitions are given below.

Definition 3.5. The three gluing moves on an abstract path $\tau = (T, \tau, \sigma)$ with $T = (V, E)$ are:

Ia. INTERIOR: Given an arrow $e \in E$ between interior nodes, contract e and set

$$\tau(v^*) = \tau(s(e)) + \tau(t(e)) - (0, (1, 0, 0), (1, 0, 0)).$$




FIGURE 7. The three gluing moves are Ia (interior), Ib (boundary-obstructed), and II (boundary). Move Ia glues two adjacent interior edges to obtain a single interior edge. Move Ib is applicable in the presence of a boundary edge such that the subpath obtained by appending adjacent interior edges is legal. Move Ib replaces this subpath with a single interior edge. Move II glues two adjacent boundary edges to obtain a single boundary edge. The effect of each gluing move is to contract the underlying tree, as shown in the second row. Interior (resp., boundary) breaks and abstract edges above correspond to black (resp., red) arrows and nodes below. A red node with a black border is valid for Move Ib (a red node without a black border may or may not be valid).

Ib. BOUNDARY-OBSTRUCTED: Given a boundary node $v \in V$ with

$$\begin{aligned} o_1(v) + s_1(v) &= |\{e \in E \mid t(e) = v\}|, \\ o_2(v) + u_2(v) &= |\{e \in E \mid s(e) = v\}|, \end{aligned}$$

contract $\text{Star}^o(v)$ and set

$$\tau(v^*) = \sum_{v' \in \text{Nbhd}^o(v)} \tau(v') - \tau(v) + (1, (0, 0, 0), (0, 0, 0)).$$

II. BOUNDARY: Given an arrow $e \in E$ between boundary nodes, contract e and set

$$\tau(v^*) = \begin{cases} \tau(s(e)) + \tau(t(e)) - (1, (0, 1, 0), (0, 1, 0)) & \text{if } \sigma(e) = u, \\ \tau(s(e)) + \tau(t(e)) - (1, (0, 0, 1), (0, 0, 1)) & \text{if } \sigma(e) = s. \end{cases}$$

We denote the resulting abstract paths by $q^{Ia}(\tau, e)$, $q^{Ib}(\tau, v)$, and $q^{II}(\tau, e)$, respectively.

Definition 3.6. A *marker* on τ is a node or arrow of the underlying tree T . We will use X as shorthand for the marker set $V \cup E$. Note that at most one of $q^{Ia}(\tau, x)$, $q^{Ib}(\tau, x)$, and $q^{II}(\tau, x)$ is defined for each $x \in X$. If one is indeed defined, we denote the resulting glued abstract path by $\tilde{q}(\tau, x)$.

3.2. Abstract paths of low weight. We now define a notion of weight for abstract paths. Since all vertices, edges, and paths in this section are abstract, we will drop the descriptor.

Definition 3.7. The *weight* of an edge $\varepsilon = (\mu, (o_1, u_1, s_1), (o_2, u_2, s_2))$ is given by:

$$w(\varepsilon) = \begin{cases} 1 & \text{if } \mu = 0, \\ 2 - s_1 - u_2 & \text{if } \mu = 1. \end{cases}$$

The *weight* of a path is the sum of the weights of its edges:

$$w(\tau) = \sum_{v \in V} w(\tau(v)).$$

Lemma 3.8. *Each gluing move reduces weight by 1.*

Proof. For Move Ia, the result is immediate.

For Move Ib, it suffices to show that the subpath τ' corresponding to $\text{Star}^o(v)$ has weight 2, since this subpath is being replaced by an edge of weight 1. We induct on the length $l(\tau')$. If $l(\tau') = 1$, then τ' consists of a single legal boundary edge $\varepsilon = \tau(v)$, which has weight 2 by definition. If $l(\tau') > 1$, then τ' includes an interior edge ε' . We suppose that ε' precedes $\varepsilon = \tau(v)$. Then the break e between ε' and ε is stable. Let τ'' be the subpath of τ' that results from pruning ε' from τ' . Then

$$w(\tau') = w(\tau'') + 1.$$

The subpath τ'' is illegal due to a single incoming stable end. Converting this stable end to an unstable end yields a legal path τ''' with

$$w(\tau''') = w(\tau'') + 1.$$

Furthermore, τ''' is a candidate for Move II of length $l(\tau') - 1$, so by induction we have

$$w(\tau''') = 2.$$

Together these equations imply $w(\tau') = 2$. The proof is similar if ε' follows ε .

For Move II, let $\varepsilon_1 = \tau(s(e))$ and $\varepsilon_2 = \tau(t(e))$ be the consecutive boundary edges and let ε be the resulting glued boundary edge. If e is stable, then

$$\begin{aligned} w(\varepsilon_1) + w(\varepsilon_2) &= (2 - s_1(\varepsilon_1) - u_2(\varepsilon_1)) + (2 - s_1(\varepsilon_2) - u_2(\varepsilon_2)) \\ &= 1 + 2 - (s_1(\varepsilon_1) + s_1(\varepsilon_2) - 1) - (u_2(\varepsilon_1) + u_2(\varepsilon_2)) \\ &= 1 + 2 + s_1(\varepsilon) - u_2(\varepsilon) \\ &= 1 + w(\varepsilon). \end{aligned}$$

The proof is similar if e is unstable. □

In order to classify paths of low weight, we introduce several counting functions.

Definition 3.9. Given a path τ , let $O(\tau)$, $U(\tau)$, and $S(\tau)$ denote the number of breaks of each type:

$$\begin{aligned} O(\tau) &= |\sigma^{-1}(o)| \\ U(\tau) &= |\sigma^{-1}(u)| \\ S(\tau) &= |\sigma^{-1}(s)| \end{aligned}$$

Let $B(\tau)$ denote the number of boundary edges:

$$B(\tau) = |\mu^{-1}(1)|$$

Proposition 3.10. *For any legal path τ , we have*

$$(10) \quad w(\tau) = 1 + O(\tau) + B(\tau).$$

In particular, a legal path has weight 1 if and only if it consists of a single interior edge.

Proof. We induct on the length $l(\tau)$. The base case $l(\tau) = 1$ follows from the fact that a legal boundary edge has weight 2.

We now suppose that $l(\tau) > 1$. Note that Move Ib, applied to a legal boundary edge, reduces both w and B by 1 while preserving O (and l), and therefore preserves (10). So we may assume without loss of generality that τ has no legal boundary edges. Now τ falls into at least one of the following cases:

Case 1: τ has two adjacent interior edges.

Case 2: τ has a boundary edge without adjacent boundary edges.

Case 3: τ has two adjacent boundary edges.

In Case 1, we can apply Move Ia. This reduces l and O by 1 and preserves B . In Case 2, since τ is legal, we can apply Move Ib to the interior star of this (illegal) boundary edge to reduce l by at least 1. This preserves O and reduces B by 1. In Case 3, we can apply Move II. This reduces l and B by 1 and preserves O . So in all three cases, the right-hand side of (10) is reduced by 1. By Lemma 3.8, the left-hand side of (10) is reduced by 1 as well. Thus (10) is preserved. Since length is reduced in all cases, we are done by induction.

We now turn to the final claim. If τ is legal and $w(\tau) = 1$, then (10) implies that τ has no boundary edges and no interior breaks. Since breaks between adjacent interior edges are necessarily interior, we conclude that τ consists of a single interior edge. □

Corollary 3.11 (Classification of legal paths of low weight).

- (i) A legal path has weight 2 if and only if it has the form Ia or Ib.
- (ii) A boundary path has length 2 if and only if it has the form II.

Proof. For (i), suppose τ is legal and $w(\tau) = 2$. By (10), there are two possibilities:

Case 1: $O(\tau) = 1$ and $B(\tau) = 0$.

Case 2: $O(\tau) = 0$ and $B(\tau) = 1$.

In Case 1, τ consists of two interior edges, giving the form Ia.

In Case 2, no two interior edges may be adjacent, so τ consists of a single boundary edge $\tau(v)$ and some number of adjacent interior edges. Thus τ coincides with $\text{Star}^o(v)$. Since τ is legal, it has the form Ib.

Claim (ii) is tautological. \square

Corollary 3.12. For any path τ , we have

$$(11) \quad w(\tau) = 1 + O(\tau) + B(\tau) - s_1(\tau) - u_2(\tau).$$

Proof. We build a legal path τ' from τ by concatenating the abstract edge $(1, (0, 1, 0), (0, 0, 1))$ before each incoming stable end of τ and after each outgoing unstable end of τ . This adds $s_1(\tau) + u_2(\tau)$ boundary edges of weight 2, so

$$\begin{aligned} w(\tau) &= w(\tau') - 2(s_1(\tau) + u_2(\tau)) \\ &= (1 + O(\tau') + B(\tau')) - 2(s_1(\tau) + u_2(\tau)) \\ &= 1 + O(\tau) + (B(\tau) + s_1(\tau) + u_2(\tau)) - 2(s_1(\tau) + u_2(\tau)) \\ &= 1 + O(\tau) + B(\tau) - s_1(\tau) - u_2(\tau) \end{aligned}$$

where the second equality uses Proposition 3.10. \square

3.2.1. Good tree-like paths.

Definition 3.13. A path is *tree-like* if every edge has exactly one outgoing end. For such a path τ , let $B_s(\tau)$ denote the number of boundary edges with stable outgoing end:

$$B_s(\tau) = |\{v \in V \mid \mu(v) = 1 \text{ and } s_2(v) = 1\}|.$$

Definition 3.14. A path is *good* if each end is interior or unstable. A break is *good* if it is interior or unstable. Let $G(\tau)$ denote the number of good breaks in τ :

$$G(\tau) = O(\tau) + U(\tau).$$

We will use the following corollary of Corollary 3.12 to characterize good tree-like paths of low weight.

Corollary 3.15. If τ is a legal tree-like path or a good tree-like path, then

$$(12) \quad w(\tau) = 1 + G(\tau) + B_s(\tau).$$

Proof. If τ ends in o or s , then

$$\begin{aligned} w(\tau) &= 1 + O(\tau) + B(\tau) \\ &= 1 + O(\tau) + (U(\tau) + B_s(\tau)) \\ &= 1 + (O(\tau) + U(\tau)) + B_s(\tau) \\ &= 1 + G(\tau) + B_s(\tau). \end{aligned}$$

If τ ends in u , then

$$\begin{aligned} w(\tau) &= 1 + O(\tau) + B(\tau) - 1 \\ &= 1 + O(\tau) + (U(\tau) + 1 + B_s(\tau)) - 1 \\ &= 1 + (O(\tau) + U(\tau)) + B_s(\tau) \\ &= 1 + G(\tau) + B_s(\tau). \end{aligned}$$

□

Proposition 3.16 (Classification of good tree-like paths of low weight).

- (i) A good tree-like path has weight 1 if and only if it has a form in (a) of Figure 8.
- (ii) A good tree-like path has weight 2 if and only if it has a form in (c) or (d) of Figure 8.

Proof. Let τ be a good tree-like path, with final edge ε .

For (i), if τ has weight 1 then by (12) all breaks of τ must be stable and follow interior edges. So if ε is interior, then no edges may precede ε . And if ε is boundary, then each incoming end of ε is preceded by at most one interior edge. These two cases yield the two forms in (a).

For (ii), If τ has weight 2 then by (12) all breaks but one are both stable and follow interior edges. Let e be this rogue break. Note that if ε is interior then it is necessarily preceded by an interior or unstable break. We are therefore left with six mutually exclusive cases:

- Case 1: e is interior and ε is interior.
- Case 2: e is interior and ε is boundary.
- Case 3: e is unstable and ε is interior.
- Case 4: e is unstable, ε is boundary, and $t(e)$ is interior.
- Case 5: e is unstable, ε is boundary, and $t(e)$ is boundary.
- Case 6: e is stable and ε is boundary.

Note that in Cases 3 and 4, the edge ε' preceding the unstable break e is boundary and all edges adjacent to ε' are interior. Since τ is good, the interior star of ε' is therefore legal.

In each case, we claim that τ must have a particular form. The six forms are shown sequentially from left to right in the second row of Figure 8. In each case, the line of argument is to perform a gluing move on τ , apply (i) to force the form of the resulting good tree-like path of weight 1, and then reverse the gluing move to recover the form of τ . In Cases 1 and 2, we apply Move Ia at e . In Cases 3 and 4, we apply Move Ib at the edge ε' preceding e . In Cases 5 and 6, we apply Move II at e . The six resulting paths are shown sequentially in the third row of Figure 8 (the seventh path does not play a role in this argument). Reversing the gluing move, we obtain the six forms in the second row. □

The following corollary of Proposition 3.16 is crucial for generalizing the structure equation to bigraphs.




FIGURE 8. The forms of a good tree-like weight 1 path are shown in (a), where the diamond and square symbols are members of the sets in (b) and ellipses indicate any finite number of incoming ends (including zero). The forms of a good, tree-like weight 2 path are shown in (c) and (d). Those with one good break are in (c), each admitting a unique gluing shown directly below in (e). Those with no good break are shown in (d), each admitting the two gluings shown in (f). The notation $\tilde{\mathcal{G}}_m^k$ is introduced in Section 4.3.

Corollary 3.17 (The key to good breaks). *Let τ be a good tree-like path of weight 2. Then τ has at most one good break. Furthermore:*

- (i) *If τ has one good break, then τ can be glued in precisely one way.*
- (ii) *If τ has no good breaks, then τ can be glued in precisely two ways.*

Proof. The path τ has at most one good break by (12). By Proposition 3.16, τ has one of the six forms shown in (c) and (d) of Figure 8. Each form in (c) has a good break and may be glued in precisely the one way shown directly below in (e). The form in (d) has no good break and may be glued in precisely the two ways shown in (f). \square

4. BIGRAPHS

We now introduce the bigraph $\tilde{\mathcal{H}}$ associated to a higraph \mathcal{H} . The “b” stands for blow-up or boundary and “bi” signifies directedness. The bigraph $\tilde{\mathcal{H}} = (\tilde{\mathcal{V}}, \tilde{\mathcal{E}})$ covers the higraph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ in the sense that there is a 3-to-1 map $\eta : \tilde{\mathcal{V}} \rightarrow \mathcal{V}$ which induces a map on edges that preserves the size of source and target. The three lifts of a vertex in \mathcal{V} are labeled by the three elements of the abstract vertex set \mathcal{V}_{ab} , and are called interior, unstable, and stable. The $2^{|I|+|J|+1}$ lifts of an edge $(I, J) \in \mathcal{E}$ are precisely the compatible abstract edges, of which $2^{|I|+|J|}$ are interior and $2^{|I|+|J|}$ are boundary. We make this precise in the following definitions, which are followed by examples in Figures 9 and 10.

Definition 4.1. The *bigraph* $\tilde{\mathcal{H}} = (\tilde{\mathcal{V}}, \tilde{\mathcal{E}})$ associated to the higraph \mathcal{H} consists of a vertex set

$$\tilde{\mathcal{V}} = \mathcal{V} \times \{o, u, s\}$$

and an edge set

$$\tilde{\mathcal{E}} \subset \mathbb{F} \times 2^{\tilde{\mathcal{V}}} \times 2^{\tilde{\mathcal{V}}}$$

with $(\mu, \tilde{I}, \tilde{J}) \in \tilde{\mathcal{E}}$ if and only if $|\eta(\tilde{I})| = |\tilde{I}|$, $|\eta(\tilde{J})| = |\tilde{J}|$, $(\eta(\tilde{I}), \eta(\tilde{J})) \in \mathcal{E}$ and $g(\tilde{I}, \tilde{J}, \mu) \in \mathcal{E}_{ab}$.

Here $\eta : 2^{\tilde{\mathcal{V}}} \rightarrow 2^{\mathcal{V}}$ is induced by the projection $\eta : \tilde{\mathcal{V}} \rightarrow \mathcal{V}$ and

$$g : \mathbb{F} \times 2^{\tilde{\mathcal{V}}} \times 2^{\tilde{\mathcal{V}}} \rightarrow \mathbb{F} \times (\mathbb{N} \times \mathbb{N} \times \mathbb{N}) \times (\mathbb{N} \times \mathbb{N} \times \mathbb{N})$$

is defined by

$$f(\mu, \tilde{I}, \tilde{J}) = (\mu, (\tilde{o}(\tilde{I}), \tilde{u}(\tilde{I}), \tilde{s}(\tilde{I})), (\tilde{o}(\tilde{J}), \tilde{u}(\tilde{J}), \tilde{s}(\tilde{J})))$$

where $\tilde{o}(\tilde{I}) = |\{(K, o) \in \tilde{I}\}|$, $\tilde{u}(\tilde{I}) = |\{(K, u) \in \tilde{I}\}|$, and $\tilde{s}(\tilde{I}) = |\{(K, s) \in \tilde{I}\}|$.

Definition 4.2. The *boundary bigraph* $\overline{\mathcal{H}} = (\overline{\mathcal{V}}, \overline{\mathcal{E}})$ associated to the higraph \mathcal{H} is the directed hypergraph given by the boundary vertices and boundary edges of $\tilde{\mathcal{H}}$. Equivalently, $\overline{\mathcal{H}}$ has the vertex set

$$\overline{\mathcal{V}} = \mathcal{V} \times \{u, s\}$$

and the edge set

$$\overline{\mathcal{E}} \subset 2^{\overline{\mathcal{V}}} \times 2^{\overline{\mathcal{V}}}$$

with $(\bar{I}, \bar{J}) \in \overline{\mathcal{E}}$ if and only if $|\eta(\bar{I})| = |\bar{I}|$, $|\eta(\bar{J})| = |\bar{J}|$, and $(\eta(\bar{I}), \eta(\bar{J})) \in \mathcal{E}$.

Definition 4.3. A *path* $\tilde{\rho} = (T, \tilde{\rho})$ in a bigraph $\tilde{\mathcal{H}}$ consists of a non-empty, directed tree $T = (V, E)$ together with a map $\tilde{\rho} : V \rightarrow \tilde{\mathcal{E}}$ such that the following properties hold.

- (1) The map $\tilde{\rho}$ induces a well-defined map $\tilde{\rho} : E \rightarrow \tilde{\mathcal{V}}$ such that

$$\{\tilde{\rho}(e)\} = t(\tilde{\rho}(s(e))) \cap s(\tilde{\rho}(t(e))).$$

- (2) If distinct arrows e and e' have the same source or the same target, then $\tilde{\rho}(e) \neq \tilde{\rho}(e')$.

Paths ρ and ρ' are *isomorphic* if there is a digraph isomorphism $f : T \rightarrow T'$ of the associated directed trees such that $\tilde{\rho} = \tilde{\rho}' f$.

As with higraphs, paths in a bigraph are unique up to unique isomorphism and we regard isomorphic paths as the same. We may similarly define the notion of a path $\bar{\rho}$ in $\bar{\mathcal{H}}$.

Let $\tilde{\mathcal{P}}$ denote the set of paths in $\tilde{\mathcal{H}}$, and let $\bar{\mathcal{P}}$ denote the set of paths in $\bar{\mathcal{H}}$. The operations of concatenation, gluing, and refinement carry over without change. Since each path $\tilde{\rho}$ determines an abstract path $g(\tilde{\rho})$, we may pull back the three gluing moves from abstract paths to bigraph paths.

Definition 4.4. A *k-step gluing sequence* on a path $\tilde{\rho}_1$ is a *k*-tuple

$$(x_1, x_2, \dots, x_k)$$

with x_i a valid marker on $\tilde{\rho}_i$ and $\tilde{\rho}_{i+1} = \tilde{q}(\tilde{\rho}_i, x_i)$. We may use the notation

$$\tilde{\rho}_1 \xrightarrow{x_1} \tilde{\rho}_2 \xrightarrow{x_2} \dots \xrightarrow{x_k} \tilde{\rho}_{k+1}$$

to emphasize the paths along the way.

As with the higraph case, gluing sequences in a bigraph are unique up to unique isomorphism. Unlike the higraph case, the number of *k*-step gluing sequences between two fixed paths in a bigraph may not be a multiple of *k*!. Indeed, Figure 12 gives an example with exactly eight 3-step gluing sequences for the bigraph in Figure 9. However, the analogue of Lemma 2.12 does hold when *k* = 2.

Lemma 4.5. *The number of 2-step gluing sequences between any two fixed paths in a bigraph is even.*

This is a special case of Proposition 4.22, whose proof we defer to Section 4.6.2.

4.2. Path DGA.

Definition 4.6. The *path DGA* of the bigraph $\tilde{\mathcal{H}}$ over the field \mathbb{F} consists of a graded algebra $\tilde{\mathcal{A}}$ equipped with a differential $\tilde{\delta} : \tilde{\mathcal{A}} \rightarrow \tilde{\mathcal{A}}$. The algebra $\tilde{\mathcal{A}}$ has underlying vector space $\mathbb{F}\langle\tilde{\mathcal{P}}\rangle$, with the subspace $\tilde{\mathcal{A}}^k$ in grading *k* spanned by the set $\tilde{\mathcal{P}}^k$ of paths of weight *k*. Multiplication is given by concatenation as before.

The value of $\tilde{\delta}$ on an interior edge $\tilde{\varepsilon}$ is the sum of all refinements of $\tilde{\varepsilon}$ to a path of weight 2. The value of $\tilde{\delta}$ on a boundary edge $\bar{\varepsilon}$ is the sum of all refinements of $\bar{\varepsilon}$ to a boundary path of length 2.

The differential $\tilde{\delta}$ has degree 1, since a boundary refinement of a boundary edge $\bar{\varepsilon}$ has length 2 if and only if it has weight $w(\bar{\varepsilon}) + 1$ by (11). We may define $\tilde{\delta}$ directly on a path $\tilde{\rho} = ((V, E), \tilde{\rho})$ by

$$(13) \quad \tilde{\delta}\tilde{\rho} = \sum_{v \in V^0} \sum_{\tilde{\rho}' \in \tilde{\mathcal{P}}^2} r((\tilde{\rho}, v), \tilde{\rho}') + \sum_{v \in V^1} \sum_{\tilde{\rho}' \in \tilde{\mathcal{P}}^2} r((\tilde{\rho}, v), \tilde{\rho}').$$

We now define the adjoint map $\tilde{\delta}$ with respect to the symmetric, bilinear pairing

$$\langle \cdot, \cdot \rangle : \tilde{\mathcal{A}} \times \tilde{\mathcal{A}} \rightarrow \mathbb{F}$$

with $\langle \tilde{\rho}_1, \tilde{\rho}_2 \rangle = 1$ if and only if $\tilde{\rho}_1 = \tilde{\rho}_2$.

Definition 4.7. The map $\tilde{\delta} : \tilde{\mathcal{A}} \rightarrow \tilde{\mathcal{A}}$ is defined on the path $\tilde{\rho}$ with underlying tree $T = (V, E)$ by

$$(14) \quad \tilde{\delta}\tilde{\rho} = \sum_{x \in X} \tilde{q}(\tilde{\rho}, x)$$

where $\tilde{q}(\tau, x)$ is defined to be zero on invalid input.

Theorem 4.8. *We have*

$$\langle \tilde{\delta}a, b \rangle = \langle a, \tilde{\delta}b \rangle$$

for all $a, b \in \tilde{\mathcal{A}}$. Furthermore, $\tilde{\partial}^2 = 0$ and $\tilde{\delta}^2 = 0$.

Proof. For the first statement, by linearity it suffices to verify that

$$\langle \tilde{\delta}\tilde{\rho}_1, \tilde{\rho}_2 \rangle = \langle \tilde{\rho}_1, \tilde{\delta}\tilde{\rho}_2 \rangle$$

for all paths $\tilde{\rho}_1$ and $\tilde{\rho}_2$. The left-hand side counts all 1-step gluing sequences $\tilde{\rho}_1 \xrightarrow{x_1} \tilde{\rho}_2$. The right-hand side counts all ways to refine an interior edge of $\tilde{\rho}_2$ to a path of weight 2 to obtain $\tilde{\rho}_1$, plus all ways to refine a boundary edge of $\tilde{\rho}_2$ to a boundary path of length 2 to obtain $\tilde{\rho}_1$. These counts coincide by Corollary 3.11.

To prove $\tilde{\partial}^2 = 0$, it suffices to show that $\langle \tilde{\partial}^2\tilde{\rho}_1, \tilde{\rho}_3 \rangle$ vanishes for all paths $\tilde{\rho}_1$ and $\tilde{\rho}_3$. Indeed, the term $\langle \tilde{\partial}^2\tilde{\rho}_1, \tilde{\rho}_3 \rangle$ counts the number of 2-step gluing sequences from $\tilde{\rho}_1$ to $\tilde{\rho}_3$, which is even by Lemma 4.5. Now $\tilde{\delta}^2 = 0$ by adjointness as before. \square

4.3. Tree-like higraphs and bigraphs. The following notions for abstract paths readily carry over to bigraphs as well.

Definition 4.9. A higraph or bigraph is *tree-like* if the target of each edge consists of exactly one vertex. A path in a bigraph is *good* if each end is interior or unstable. A break is *good* if it is interior or unstable.

Clearly the bigraph $\tilde{\mathcal{H}}$ is tree-like if and only if the underlying higraph \mathcal{H} is tree-like. Let $\tilde{\mathcal{G}}^k$ be the set of good paths in $\tilde{\mathcal{H}}$ of weight k . Let $\tilde{\mathcal{G}}_m^k$ be the set of good paths in $\tilde{\mathcal{H}}$ of weight k with m good breaks. Note that $\tilde{\mathcal{G}}^1 = \tilde{\mathcal{G}}_0^1$ and $\tilde{\mathcal{G}}^2 = \tilde{\mathcal{G}}_1^2 \cup \tilde{\mathcal{G}}_0^2$ by (12); the forms of these paths are shown in Figure 8.

Definition 4.10. We define the element $\tilde{D} \in \tilde{\mathcal{A}}$ by

$$\tilde{D} = \sum_{\tilde{\rho} \in \tilde{\mathcal{G}}^1} \tilde{\rho}.$$

Theorem 4.11. *If \mathcal{H} is tree-like, then the element \tilde{D} satisfies the structure equation:*

$$\tilde{\delta}\tilde{D} = \tilde{D} \circ \tilde{D}.$$

Proof. We claim

$$(15) \quad \tilde{\delta}\tilde{D} = \sum_{\tilde{\rho} \in \tilde{\mathcal{G}}_1^2} \tilde{\rho} + \sum_{\tilde{\rho} \in \tilde{\mathcal{G}}_0^2} 2\tilde{\rho} = \sum_{\tilde{\rho} \in \tilde{\mathcal{G}}_1^2} \tilde{\rho} = \tilde{D} \circ \tilde{D}.$$

The left-hand inequality holds by Corollary 3.17. The middle equality holds over \mathbb{F} . The right-hand equality holds because each good path of weight 2 with one good break arises as a product of two (not necessarily distinct) good paths of weight 1 in a unique way (as specified by cutting at the good break). \square

4.4. Balanced higraphs and bigraphs. The reader may check that Theorem 4.11 is always false if \mathcal{H} is not tree-like. On the other hand, there is a duality at play, which switches the roles of unstable and stable vertices. Dual to tree-like higraphs are *rake-like* higraphs, where the source of each edge consists of a single vertex. If \mathcal{H} is rake-like, and we re-define good breaks to be interior or stable ones, then Theorem 4.11 continues to hold. More generally, we can independently define the notion of good at each vertex of \mathcal{H} , provided each edge of \mathcal{H} has a compatible shape.

Definition 4.12. A *vertex assignment* on a higraph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ is a map $\psi : \mathcal{V} \rightarrow \{1, 2\}$ such that

$$(16) \quad |s(\varepsilon) \cap \psi^{-1}(2)| + |t(\varepsilon) \cap \psi^{-1}(1)| = 1$$

for every edge $\varepsilon \in \mathcal{E}$. In this case, we say that \mathcal{H} is *balanced* by ψ . The set of ψ -good vertices of $\tilde{\mathcal{H}}$ is given by

$$\tilde{\mathcal{V}}_\psi = (\psi^{-1}(1) \times \{o, u\}) \cup (\psi^{-1}(2) \times \{o, s\}) \subset \tilde{\mathcal{V}}.$$

A path $\tilde{\rho}$ is ψ -good if each end is ψ -good. A break e of $\tilde{\rho}$ is ψ -good if $\tilde{\rho}(e)$ is ψ -good. Let $\tilde{\mathcal{P}}^k(\psi)$ denote the set of ψ -good paths of weight k . Let $\tilde{\mathcal{P}}_m^k(\psi)$ denote the set of ψ -good paths of weight k with m ψ -good breaks.

Note if $\psi \equiv 1$ then $\tilde{\mathcal{P}}_m^k(\psi) = \tilde{\mathcal{G}}_m^k$, and ψ balances \mathcal{H} if and only if \mathcal{H} is tree-like. If $\psi \equiv 2$, then ψ -good breaks are interior or stable ones, and ψ balances \mathcal{H} if and only if \mathcal{H} is rake-like.

Remark 4.13. The proof of the monopole surgery exact triangle is modeled by a higraph with $\psi = 1$ for three vertices and $\psi = 2$ for four vertices, as explained in Figure 19.

Definition 4.14. We define the element $\tilde{D}_\psi \in \tilde{\mathcal{A}}$ by

$$\tilde{D}_\psi = \sum_{\tilde{\rho} \in \tilde{\mathcal{P}}^1(\psi)} \tilde{\rho}.$$

Corollary 4.15. *If \mathcal{H} is balanced by ψ , then the element \tilde{D} satisfies the structure equation:*

$$\tilde{\delta}\tilde{D}_\psi = \tilde{D}_\psi \circ \tilde{D}_\psi.$$

Proof. Given $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ balanced by ψ , we set $\mathcal{H}' = (\mathcal{V}, \mathcal{E}')$ with \mathcal{E} in bijection with \mathcal{E}' via

$$\varepsilon = (I, J) \mapsto \varepsilon' = ((I \cap \phi^{-1}(1)) \cup (J \cap \phi^{-1}(2)), (J \cap \phi^{-1}(1)) \cup (I \cap \phi^{-1}(2))).$$

In other words, we flip each end with $\psi = 2$ from source to target or vice versa. The equation

$$\tilde{\delta}\tilde{D} = \sum_{\tilde{\rho} \in \tilde{\mathcal{P}}_1^2(\psi)} \tilde{\rho} + \sum_{\tilde{\rho} \in \tilde{\mathcal{P}}_0^2(\psi)} 2\tilde{\rho} = \sum_{\tilde{\rho} \in \tilde{\mathcal{P}}_1^2(\psi)} \tilde{\rho} = \tilde{D} \circ \tilde{D}$$

on $\tilde{\mathcal{H}}$ is equivalent to equation (15) on $\tilde{\mathcal{H}}'$. The latter holds since \mathcal{H}' is tree-like by (16). \square

4.5. Boundary bigraphs. We briefly return to boundary bigraphs.

Definition 4.16. The *path DGA* of the boundary bigraph $\overline{\mathcal{H}}$ over the field \mathbb{F} consists of a graded algebra $\overline{\mathcal{A}}$ equipped with a differential $\bar{\delta} : \overline{\mathcal{A}} \rightarrow \overline{\mathcal{A}}$. The algebra $\overline{\mathcal{A}}$ has underlying vector space $\mathbb{F}\langle\overline{\mathcal{P}}\rangle$, with the subspace $\overline{\mathcal{A}}^k$ in grading k spanned by the set $\overline{\mathcal{P}}^k$ of paths of length k . Multiplication is given by concatenation.

The value of $\bar{\delta}$ on a boundary edge $\bar{\varepsilon}$ equals the sum of all refinements of $\bar{\varepsilon}$ to a boundary path of length 2. In particular, $\overline{\mathcal{A}}$ sits inside $\tilde{\mathcal{A}}$ as a differential algebra.

The arguments in Section 2 for higraphs carry over essentially without change to boundary bigraphs. In particular, $\bar{\delta}$ is indeed a differential and the element

$$\bar{D} = \sum_{\bar{\rho} \in \overline{\mathcal{P}}^1} \bar{\rho}$$

of $\overline{\mathcal{A}}$ satisfies the structure equation

$$\bar{\delta}\bar{D} = \bar{D} \circ \bar{D}$$

without restriction on the shape of edges.

4.6. Gluing sequences. We now take a closer look at gluing sequences in order to prove Lemma 2.12 and Lemma 4.5. A path in a higraph with two breaks has a 2-step gluing sequence. However, there exist paths in a bigraph, such as path (d) in Figure 8, with two available gluing moves but no 2-step gluing sequence. The subtlety is that the breaks in a gluing sequence live on different paths. At least for bigraphs, it is not immediately clear how to go about reordering them. The key will be to pull the breaks or markers in a gluing sequence up to breaks on the initial path. Though the higraph case is straightforward, we consider it in detail as a model for the bigraph case.

4.6.1. *Higraphs.* Let ρ_1 be a path in a higraph \mathcal{H} . We will define a bijection between k -tuples of distinct breaks on ρ_1 and k -step gluing sequences on ρ_1 . We relate these notions by populating a lower triangular matrix A whose (i, j) -entry e_j^i is a break on ρ_j , with $\rho_j = q(\rho_{j-1}, e_{j-1}^{j-1})$. Here we use the bijection π on arrows as defined following (8).

Let $(e_1^1, e_2^1, \dots, e_k^1)$ be a k -tuple of distinct breaks on a path ρ_1 , thought of as the first column of A . We populate A one column at a time from left to right as follows. For each j from 2 to k , we set

$$\rho_j = q(\rho_{j-1}, e_{j-1}^{j-1})$$

and

$$e_j^i = \pi(e_{j-1}^i)$$

for $j \leq i \leq k$. The diagonal of A is a k -step gluing sequence $(e_1^1, e_2^2, \dots, e_k^k)$ on ρ_1 .

Let $(e_1^1, e_2^2, \dots, e_k^k)$ be a k -step gluing sequence on ρ_1 , thought of as the diagonal of A . We populate A one column at a time from right to left as follows. For each j from k down to 2, we have

$$\rho_j = q(\rho_{j-1}, e_{j-1}^{j-1})$$

and can thus set

$$e_{j-1}^i = \pi^{-1}(e_j^i)$$

for $j \leq i \leq k$. The first column of A is a k -tuple $(e_1^1, e_2^1, \dots, e_k^1)$ of distinct breaks on ρ_1 .

Proposition 4.17. *Let ρ_1 be a path in a higraph. The map*

$$(e_1^1, e_2^1, \dots, e_k^1) \mapsto (e_1^1, e_2^2, \dots, e_k^k)$$

gives a bijection between k -tuples of distinct breaks on ρ_1 and k -step gluing sequences on ρ_1 .

Furthermore, if two k -tuples differ by a permutation, then the corresponding k -step gluing sequences terminate in the same path ρ_{k+1} .

Proof. The map is a bijection as we have also defined its inverse above.

For the latter statement, consider two k -tuples $(e_1^1, e_2^1, e_3^1, \dots, e_k^1)$ and $(f_1^1, f_2^1, f_3^1, \dots, f_k^1)$ of distinct breaks on ρ_1 which differ by a permutation, which we may assume is a transposition of neighbors. We may reduce further to the case of the transposition of the first two entries by induction on k , since if $e_1^1 = f_1^1$ then (e_2^1, \dots, e_k^1) and (f_2^1, \dots, f_k^1) are $(k-1)$ -tuples of distinct breaks on $\rho_2 = q(\rho_1, e_1^1) = q(\rho_1, f_1^1)$. Note that the terminal path of (e_2^1, \dots, e_k^1) is the same as that of $(e_1^1, e_2^1, e_3^1, \dots, e_k^1)$ by definition, and similarly for (f_2^1, \dots, f_k^1) .

So we suppose $e_1^1 = f_2^1$, $e_2^1 = f_1^1$, and $e_i^1 = f_i^1$ for $3 \leq i \leq k$. Let T_3^e and T_3^f be the trees underlying the paths $q(q(\rho_1, e_1^1), e_2^1)$ and $q(q(\rho_1, f_1^1), f_2^1)$. When e_1^1 and e_2^1 are adjacent, these trees result from contracting T_1 along the same two-arrow subtree T' , and we therefore have an isomorphism $f : T_3^e \rightarrow T_3^f$ which identifies $\pi_{e_2^1}(\pi_{e_1^1}(v))$ with $\pi_{f_2^1}(\pi_{f_1^1}(v))$ for each $v \in V_1$, and identifies the arrow e_3^1 with f_3^1 for $3 \leq i \leq k$. Furthermore, the map f is an isomorphism of paths, since for both paths the edge over the node v^* to which T' contracts has the same source and target as the subpath of ρ_1 supported by T' . Finally, if e_1^1 and e_2^1 are not adjacent, then the isomorphism of paths is clear.

Therefore the paths in the gluing sequences $(e_1^1, e_2^1, e_3^1, \dots, e_k^1)$ and $(f_1^1, f_2^1, f_3^1, \dots, f_k^1)$ are identified from ρ_3 to ρ_{k+1} . \square

Proof of Lemma 2.12. Via Proposition 4.17, the set of k -step gluing sequences between any two fixed paths admits a free action of the group of permutations on k elements. \square

4.6.2. Bigraphs. We now take an analogous, though more involved, approach to bigraphs. We will define a bijection between certain k -tuples of markers on ρ_1 and k -step gluing sequences on ρ_1 . Recall the notion of a marker from Definition 3.6.

Definition 4.18. A marker x on $\tilde{\rho}_1$ is *valid* if $\tilde{q}(\tilde{\rho}_1, x)$ is defined. A k -tuple (x_1^1, \dots, x_1^k) of markers on $\tilde{\rho}_1$ is *valid* if its coordinates are valid and distinct.

Suppose x_1^1 is a valid marker on $\tilde{\rho}_1$, and set $\tilde{\rho}_2 = \tilde{q}(\tilde{\rho}_1, x_1^1)$. Let $T' = (V', E')$ be the subtree of T_1 that is contracted under the gluing move, and let v^* be the node in T_2 onto which T' contracts. Recall that π denotes both the map $\pi : V_1 \rightarrow V_2$ which is injective on $V_1 - V'$ with image $V_2 - \{v^*\}$, and the bijection $\pi : E_1 - E' \rightarrow E_2$. We thus have maps $\pi : X_1 - E' \rightarrow X_2$ and $\pi^{-1} : X_2 - \{v^*\} \rightarrow X_1$ which are mutual inverses whenever their composition is defined.

We define maps $\tilde{\pi} : X_1 - E' \rightarrow X_2$ and $\tilde{\pi}^{-1} : X_2 \rightarrow X_1$ as follows:

$$\tilde{\pi}(x) = \begin{cases} \pi(s(x)) & \text{if } x \text{ is a boundary arrow and } t(x) = x_1^1, \\ \pi(t(x)) & \text{if } x \text{ is a boundary arrow and } s(x) = x_1^1, \\ \pi(x) & \text{otherwise.} \end{cases} \quad \begin{array}{l} (1a) \\ (1b) \\ (1c) \end{array}$$

$$\tilde{\pi}^{-1}(y) = \begin{cases} (\pi^{-1}(y), x_1^1) & \text{if } y \text{ is a boundary node and } (y, \pi(x_1^1)) \in E_2, \\ (x_1^1, \pi^{-1}(y)) & \text{if } y \text{ is a boundary node and } (\pi(x_1^1), y) \in E_2, \\ x_1^1 & \text{if } y = \pi(x_1^1), \\ t(x_1^1) & \text{if } y = \pi(t(x_1^1)) \text{ and } x_1^1 \text{ is unstable,} \\ s(x_1^1) & \text{if } y = \pi(s(x_1^1)) \text{ and } x_1^1 \text{ is stable or interior,} \\ \pi^{-1}(y) & \text{otherwise.} \end{cases} \quad \begin{array}{l} (2a) \\ (2b) \\ (2c) \\ (2d) \\ (2e) \\ (2f) \end{array}$$

Note that in (1a) and (1b), x_1^1 is a marker for Move Ib. In (2a), (2b), and (2c), x_1^1 is a marker for Move Ib. In (2d), x_1^1 is a marker for Move II, while in (2e), x_1^1 is a marker for Move II or Move Ia. In (2c), (2d), and (2e), y coincides with the node v^* .

The following facts will be useful in the proof of Lemma 4.19 below:

Fact 1: The maps $\pi|_{X_1 - V' - E'}$ and $\pi^{-1}|_{X_2 - \{v^*\}}$ are mutual inverses and satisfy $\tilde{\rho}_2\pi = \tilde{\rho}_1$ and $\tilde{\rho}_2 = \tilde{\rho}_1\pi^{-1}$. In particular, these maps preserve abstract vertex and edge type.

Fact 2: If x_1^1 is a valid marker for Move Ib, then x_1^1 is boundary while $\pi(x_1^1)$ is interior and thus not valid.

Lemma 4.19. *The maps $\tilde{\pi} : X_1 - E' \rightarrow X_2$ and $\tilde{\pi}^{-1} : X_2 \rightarrow X_1$ are well-defined. Furthermore:*

- (i) *If $y \in X_2$, then $\tilde{\pi}(\tilde{\pi}^{-1}(y)) = y$. If y is valid, then $\tilde{\pi}^{-1}(y)$ is distinct from x_1^1 and valid.*
- (ii) *If $x \in X_1$ is distinct from x_1^1 and valid, then $\tilde{\pi}(x)$ is defined and $\tilde{\pi}^{-1}(\tilde{\pi}(x)) = x$.*

Proof. The cases within each definition are mutually exclusive. The map $\tilde{\pi}$ is well-defined because its domain coincides with that of π . For the map $\tilde{\pi}^{-1}$, there are two issues to check. First, Fact 1 implies that $(\pi^{-1}(y), x_1^1)$ is an arrow of T_1 in (2a) and that $(x_1^1, \pi^{-1}(y))$ is an arrow of T_1 in (2b). Second, if $y = v^*$, then $\tilde{\pi}^{-1}(y)$ is defined by (2c) if x_1^1 is a node and by (2d) or (2e) when x_1^1 is an arrow.

For the first claim of (i), note that the image of $\tilde{\pi}^{-1}$ is disjoint from E' , so that $\tilde{\pi}(\tilde{\pi}^{-1}(y))$ is well-defined. Cases (1a) and (2a) define mutual inverses, as do (1b) and (2b). Case (1c) inverts (2c), (2d), (2e), and (2f) by definition (thought not necessarily the other way around).

Now suppose that $y \in X_2$ is valid. If x_1^1 is a node, then $v^* = \pi(x_1^1)$ is an interior node and thus not valid. Nor is x_1^1 in the image of π^{-1} . So $\tilde{\pi}^{-1}(y)$ is distinct from x_1^1 .

To show that $\tilde{\pi}^{-1}(y)$ is valid, we consider each case of the definition of $\tilde{\pi}^{-1}$ in turn.

Case (2a): y is a boundary node with $(y, \pi(x_1^1)) \in E_2$. Then $\tilde{\pi}^{-1}(y)$ is the arrow $(\pi^{-1}(y), x_1^1)$ which runs between two boundary nodes by Facts 1 and 2 and is thus valid.

Case (2b): This case is similar (2a).

Case (2c): This case does not occur when y is valid by Fact 2.

Case (2d): $y = v^*$ and x_1^1 is an unstable arrow, with $\tilde{\pi}^{-1}(y) = t(x_1^1)$. Since x_1^1 is a valid boundary arrow, $t(x_1^1)$ is a boundary node. The edge $\tilde{\rho}_2(y)$ results from applying Move II to the subpath $\tilde{\rho}_1(s(x_1^1))\tilde{\rho}_1(t(x_1^1))$. Furthermore, by Fact 2, the interior neighborhoods of y and $t(x_1^1)$ are identified. Since y is valid and the break $\tilde{\rho}_1(x_1^1)$ is unstable, we conclude that $t(x_1^1)$ is valid (while $s(x_1^1)$ is not).

Case (2e): This case is similar to Case 4 when x_1^1 is stable. Note that if $y = v^*$ is valid, then x_1^1 cannot be an interior arrow.

Case (2f): If y is a valid arrow then its ends are both interior or both boundary, and it suffices to show that the same is true of $\tilde{\pi}^{-1}(y) = \pi^{-1}(y)$. This follows from Fact 1 if y is not bounded by the interior node v^* . On the other hand, if y starts or ends at v^* , then both ends of y are interior nodes, and thus y is an interior arrow. Hence $\pi^{-1}(y)$ is also an interior arrow by Fact 1, and we conclude that its ends are both interior as well. If y is a valid node and does not fall into (2a) through (2e), then $\text{Star}^o(y)$ does not include v^* and thus $\text{Star}^o(\pi^{-1}(y))$ is legal as well by Fact 1. Therefore $\pi^{-1}(y)$ is also a valid node. This completes (i).

For (ii), we suppose $x \in X_1$ is distinct from x_1^1 and valid. If x_1^1 is an arrow, then $E' = \{x_1^1\}$ so $x \notin E'$. If x_1^1 is a node, then E' consists of arrows between the boundary node x_1^1 and interior nodes. These arrows are not valid, so $x \notin E'$. Therefore $\tilde{\pi}(x)$ is defined.

Finally, we show that $\tilde{\pi}^{-1}(\tilde{\pi}(x)) = x$. As in (i), cases (1a) and (2a) are mutual inverses, as are (1b) and (2b). So from now on we assume x falls into (1c), with $\tilde{\pi}(x) = \pi(x)$. If $\pi(x) = \pi(x_1^1)$ as in (2c), then $x = x_1^1$ since the nodes adjacent to x_1^1 are not valid. If $\pi(x)$ falls into (2d) or (2e), then x_1^1 is a boundary arrow bounded by the valid node x and $\pi(x) = \pi(s(x_1^1)) = \pi(t(x_1^1))$. If x_1^1 is unstable (resp., stable) then we must have $x = t(x_1^1)$ (resp., $x = s(x_1^1)$) in order for $\text{Star}^o(x)$ to be legal. This is consistent with the definition of $\tilde{\pi}^{-1}$ in (1d) and (1e). Finally, if $\pi(x)$ falls into (1f), then $\tilde{\pi}^{-1}(\tilde{\pi}(x)) = \pi^{-1}(\pi(x)) = x$ by Fact 1. \square

Paralleling our strategy for higraph, we now populate a lower triangular matrix whose (i, j) -entry x_j^i is a valid marker on $\tilde{\rho}_j$, with $\tilde{\rho}_j = \tilde{q}(\tilde{\rho}_{j-1}, x_{j-1}^{j-1})$.

Let $(x_1^1, x_1^2, \dots, x_1^k)$ be a k -tuple of markers on a path $\tilde{\rho}_1$. For each j from 2 to k , we set

$$\tilde{\rho}_j = \tilde{q}(\tilde{\rho}_{j-1}, x_{j-1}^{j-1})$$

and

$$x_j^i = \tilde{\pi}(x_{j-1}^i)$$

for $j \leq i \leq k$. This inductive definition may break down if some x_j^i is not a valid marker.

Definition 4.20. A k -tuple $(x_1^1, x_1^2, \dots, x_1^k)$ of valid markers on $\tilde{\rho}_1$ is *completely valid* if x_j^i is defined and valid for all $1 \leq j \leq i \leq k$.

Thus, a completely valid k -tuple $(x_1^1, x_1^2, \dots, x_1^k)$ on $\tilde{\rho}_1$ determines a k -step gluing sequence $(x_1^1, x_2^2, \dots, x_k^k)$ on $\tilde{\rho}_1$.

Let $(x_1^1, x_2^2, \dots, x_k^k)$ be a k -step gluing sequence on $\tilde{\rho}_1$. For each j from k down to 2, we have

$$\tilde{\rho}_j = \tilde{q}(\tilde{\rho}_{j-1}, x_{j-1}^{j-1})$$

and can thus set

$$x_{j-1}^i = \tilde{\pi}^{-1}(x_j^i)$$

for $j \leq i \leq k$. This determines a completely valid k -tuple $(x_1^1, x_1^2, \dots, x_1^k)$ of markers on $\tilde{\rho}_1$ by Lemma 4.19, which also implies that the above maps are inverses. We conclude:

Proposition 4.21. *Let $\tilde{\rho}_1$ be a path in a bigraph. The map*

$$(x_1^1, \dots, x_1^k) \mapsto (x_1^1, \dots, x_k^k)$$

gives a bijection between completely valid k -tuples of markers on $\tilde{\rho}_1$ and k -step gluing sequences on $\tilde{\rho}_1$.

In a bigraph, permutations of a completely valid k -tuple need not be completely valid. Indeed, in Figure 12 each set of 3 valid markers on $\tilde{\rho}_1$ is completely valid for only two out of six possible orderings, and there are precisely eight 3-step gluing sequences. However, Lemma 2.12 does extend from higraphs to bigraphs when $k = 2$.

Proposition 4.22. *Let (x_1^1, x_1^2) be a completely valid 2-tuple for a path in a bigraph.*

- (i) *The transposed 2-tuple (x_1^2, x_1^1) is completely valid.*
- (ii) *The 2-step gluing sequences for (x_1^1, x_1^2) and (x_1^2, x_1^1) terminate in the same path.*
- (iii) *For $k \geq 2$, the number of k -step gluing sequences between any two fixed paths is even.*

Proof. For (i), Let (x_1^1, x_1^2) be a valid 2-tuple of markers on $\tilde{\rho}_1$ and let (y_1^1, y_1^2) be its transposition, so $y_1^1 = x_1^2$, $y_1^2 = x_1^1$. By definition, we have $x_2^2 = \tilde{\pi}_{x_1^1}(x_1^2)$ and $y_2^2 = \tilde{\pi}_{y_1^1}(y_1^2)$, where the subscript on $\tilde{\pi}$ specifies the marker at which we glued. We claim that the marker x_2^2 on $\tilde{q}(\tilde{\rho}_1, x_1^1)$ is valid if and only if the marker y_2^2 on $\tilde{q}(\tilde{\rho}_1, y_1^1)$ is valid.

By Fact 1, the claim is clear if x_1^1 and x_1^2 do not interact; that is, if the subtrees of T_1 that are contracted by gluing at x_1^1 and x_1^2 are disjoint. There are five ways in which two valid markers can interact. These are pictured in columns (a) through (e) of Figure 11. One can check that x_2^2 and y_2^2 are always both valid in all cases but (e). In column (e), x_1^1 is a boundary node, x_1^2 is a boundary arrow, and $t(x_1^2) = x_1^1$ or $s(x_1^2) = x_1^1$. In this case, Figures 13 and 14 show that x_2^2 is valid if and only if y_2^2 is valid.

For (ii), let T_3^x and T_3^y be the trees underlying the paths $\tilde{q}(\tilde{q}(\tilde{\rho}_1, x_1^1), x_2^2)$ and $\tilde{q}(\tilde{q}(\tilde{\rho}_1, y_1^1), y_2^2)$. In cases (a) through (e) of Figure 11, one can check that the trees result from contracting the same subtree T' of the tree T_1 , and we therefore have an isomorphism $f : T_3^x \rightarrow T_3^y$ which identifies $\pi_{x_2^2}(\pi_{x_1^1}(v))$ with $\pi_{y_2^2}(\pi_{y_1^1}(v))$ for each $v \in V_1$. The most interesting case (e) is verified in Figure 13. Furthermore, in each case, the map f is an isomorphism of paths, since for both paths the edge over the node v^* to which T' contracts has the same source and target as the subpath of ρ_1 supported by T' . Finally, if x_1^1 and x_1^2 do not interact, then the isomorphism of paths is clear.

For (iii), note that we have shown that the set of 2-step gluing sequences between any two fixed paths admits a fixed-point free involution, given by transposition on the corresponding 2-tuple of valid markers. For $k > 2$, the set of k -step gluing sequences admits a fixed-point free involution as well, since each k -step gluing sequence factors uniquely as a $(k - 2)$ -step gluing sequence followed by a 2-step gluing sequence. We conclude that these sets contain an even number of elements. \square

	(a)	(b)	(c)	(d)	(e)	(f)	(g)	(h)
x_1^1 x_2^2	y_1^1 y_2^2	Ia Ia	Ib Ib	II II	Ia Ib	Ib † II †	Ia II	Ib II
		Ia Ia	Ib Ib	II II	Ib Ia	Ib † II †	II Ia	Ib Ib
subpath of T_1								
image in T_3								

FIGURE 11. To a 2-tuple (x_1^1, x_1^2) of valid markers, we may associate the minimal undirected subpath of T_1 which includes both markers. Several such subpaths are depicted in the second row above. There are exactly five (undirected) subpaths that may arise when the contracting trees of x_1^1 and x_1^2 overlap. These are depicted in columns (a) through (e). In the first row, we depict the gluing type of the markers x_1^1 , x_2^2 , y_1^1 , and y_2^2 , where (y_1^1, y_1^2) is the transposition of (x_1^1, x_1^2) . In all columns except (e), the markers x_2^2 and y_2^2 are always valid. In column (e), x_2^2 is valid if and only if y_2^2 is valid, as shown in Figures 13 and 14. If x_2^2 and y_2^2 are valid, then $\tilde{\rho}_3$ is defined and independent of gluing order, as may be easily verified in each case above. In the third row, we show the image in T_3 of the above subpath of T_1 .




FIGURE 12. The legal path $\tilde{\rho}_1$ at left has weight 4, so all 3-step gluing sequences terminate in the path $\tilde{\rho}_4$ consisting of a single interior edge. These eight sequences are shown above, with each arrow positioned to pick out the corresponding gluing move. The four valid markers of $\tilde{\rho}_1$ are numbered on the underlying tree T_1 , with 1 and 4 marking Move Ib and 2 and 3 marking Move II. Note that the 2-tuples (1, 2), (2, 1), (3, 4) and (4, 3) fall into the undefined case of column (e) of Figure 11, as illustrated in Figure 14. The eight completely valid 3-tuples on $\tilde{\rho}_1$ are (1, 3, 4), (1, 4, 3), (2, 3, 4), (2, 4, 3), (3, 1, 2), (3, 2, 1), (4, 1, 2), and (4, 2, 1).

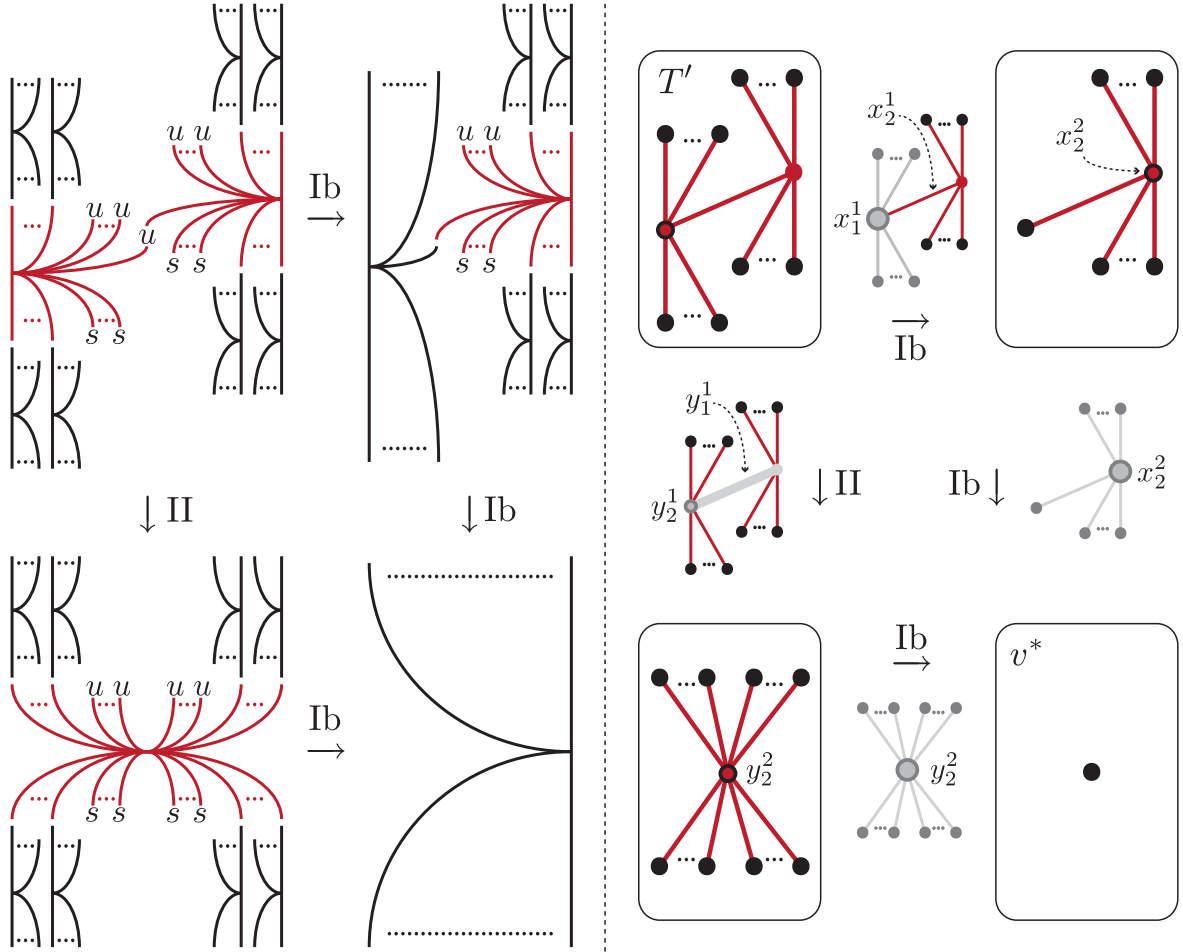


FIGURE 13. In Figures 13 and 14, we examine the case of column (e) of Figure 11 in detail. The above diagram illustrates the case when both gluing sequences (x_1^1, x_2^2) and (y_1^1, y_2^2) are defined. The subtree contracted at each step is shown in gray at right. Overall, the two sequences contract the same subtree T' of T_1 to a single node v^* , and therefore terminate in the same path $\tilde{\rho}_3$. The diagram further illustrates that x_2^2 is valid if and only if y_2^2 is valid (compare with Figure 14).

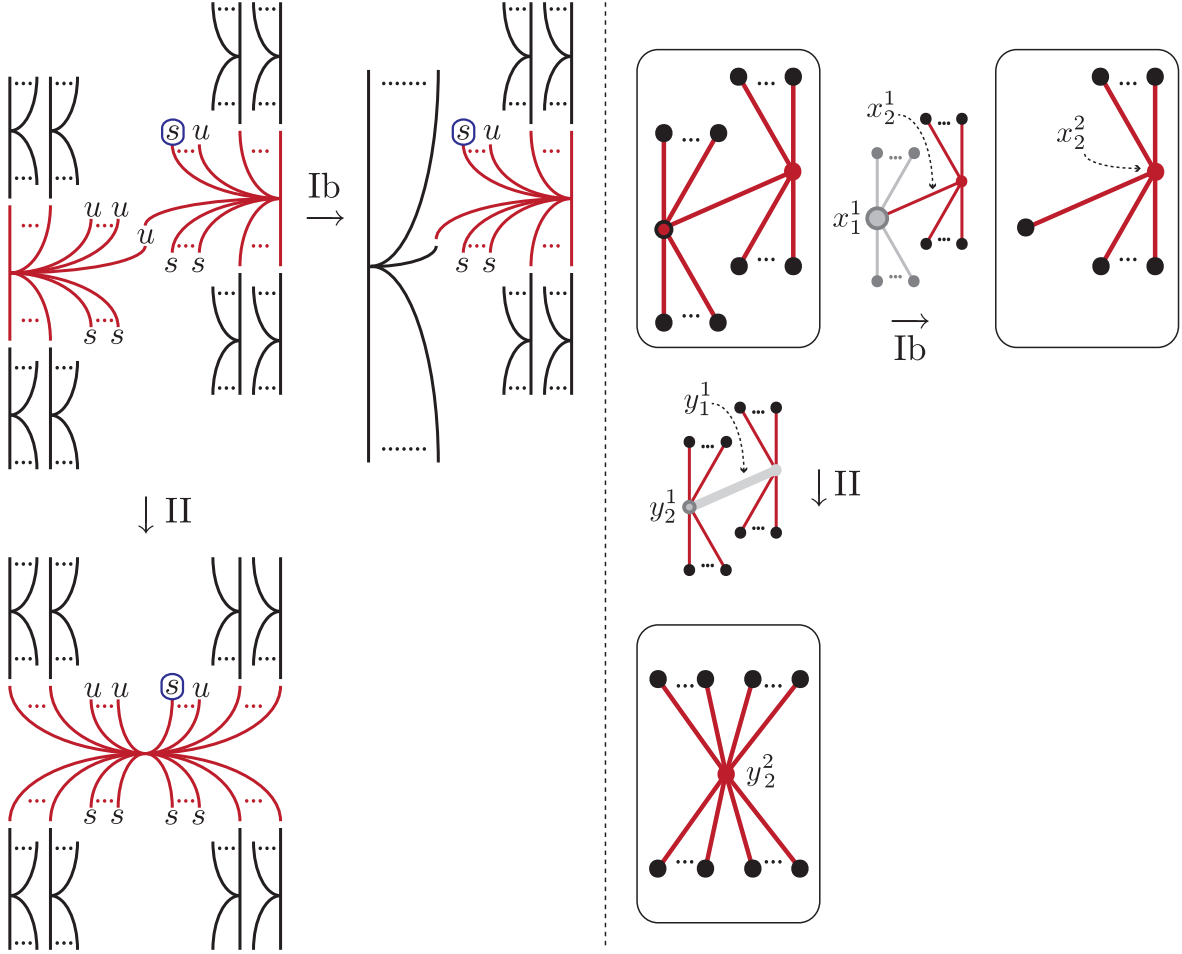


FIGURE 14. In the above diagram, both x_2^2 and y_2^2 are invalid, so that neither 2-step gluing sequence is defined. The difference between Figures 13 and 14 amounts to switching the single (encircled) vertex from unstable to stable.

APPENDIX: MORSE HOMOLOGY WITH BOUNDARY

A. Morse homology on a closed manifold. We first recall the construction of Morse homology on a closed, smooth manifold M . Fix a Riemannian metric and Morse function $f : M \rightarrow \mathbb{R}$. The gradient flow $\phi_t : M \rightarrow M$ with respect to $-\nabla f$ determines an unstable manifold

$$S_a = \{x \in M : \lim_{t \rightarrow -\infty} \phi_t(x) = a\}$$

and stable manifold

$$U_a = \{x \in M : \lim_{t \rightarrow \infty} \phi_t(x) = a\}$$

for each critical point a . The Morse index λ_a of a is the dimension of U_a . We insist that f satisfy the *Morse-Smale* condition:

- U_a and S_b intersect transversely in M for all critical points a and b .

The manifold $U_a \cap S_b$ has dimension $\lambda_a - \lambda_b$ and is canonically identified with the moduli space of gradient trajectories from a to b :

$$\mathcal{M}(a, b) = \{\gamma : \mathbb{R} \rightarrow M : \frac{d\gamma}{dt} = -\nabla f \circ \gamma, \lim_{t \rightarrow -\infty} \gamma(t) = a, \lim_{t \rightarrow +\infty} \gamma(t) = b\}.$$

The space $\check{\mathcal{M}}(a, b)$ of unparameterized gradient trajectories is the quotient of $\mathcal{M}(a, b)$ by the \mathbb{R} -translation $\gamma(\cdot) \mapsto \gamma(\cdot + s)$, so

$$(17) \quad \dim \check{\mathcal{M}}(a, b) = \lambda_a - \lambda_b - 1.$$

The Morse complex $C(M)$ is the \mathbb{F} -vector space with basis the set of critical points:

$$C(M) = \bigoplus_{a \in \text{crit}(f)} \mathbb{F}a$$

The differential ∂ counts unparameterized gradient trajectories:

$$\langle \partial a, b \rangle = |\check{\mathcal{M}}(a, b)|.$$

This notation signifies that the coefficient of b in ∂a is the number of unparameterized gradient trajectories from a to b , with the convention $|\cdot| = 0$ on infinite sets. The Morse complex is graded by Morse index and the differential has degree -1 by (17). The Morse homology $H_*(C(M), \partial)$ is isomorphic to the singular homology of M .

B. Morse homology on a manifold with boundary. Now let M be a smooth manifold with boundary. The manifold M may be identified with the quotient of a smooth double

$$\tilde{M} = M \cup_{\partial M} -M$$

by the natural involution $i : \tilde{M} \rightarrow \tilde{M}$. An i -invariant metric and Morse function on \tilde{M} determine a metric and Morse function on M by restriction. For such a metric and Morse function on M , the gradient field ∇f is tangent to ∂M along the boundary. Furthermore, at each critical point on the boundary, the normal vector is an eigenvector of the Hessian $\nabla^2 f$. A boundary critical point is called *stable* or *unstable* according to whether the corresponding eigenvalue is positive or negative, respectively. Hence we may sort the critical points of f into three types: interior, (boundary-)stable, and (boundary-)unstable. We denote these types by the letters o , s , and u . If a trajectory $\gamma \in \check{\mathcal{M}}^\partial(a, b)$ intersects the interior of M , then a must

be interior or unstable and b must be interior or stable. On the other hand, if a is stable and b is unstable, then both U_a and S_b are contained in ∂M and the most we can ask is that U_a and S_b intersect transversely as submanifolds of ∂M . Indeed, we insist that f is *regular*, which means

- U_a and S_b intersect transversely in ∂M if a is stable and b is unstable.
- U_a and S_b intersect transversely in M if a is not stable or b is not unstable.

Note that regular is equivalent to Morse-Smale if $\partial M = \emptyset$.

Next we define eight linear maps between the vector spaces $C^o(M)$, $C^s(M)$, and $C^u(M)$ generated by the interior, stable, and unstable critical points, respectively:

$$\begin{array}{ll} \partial_o^o : C^o(M) \rightarrow C^o(M) & \bar{\partial}_s^s : C^s(M) \rightarrow C^s(M) \\ \partial_s^o : C^o(M) \rightarrow C^s(M) & \bar{\partial}_u^s : C^s(M) \rightarrow C^u(M) \\ \partial_o^u : C^u(M) \rightarrow C^o(M) & \bar{\partial}_u^u : C^u(M) \rightarrow C^u(M) \\ \partial_s^u : C^u(M) \rightarrow C^s(M) & \bar{\partial}_s^u : C^u(M) \rightarrow C^s(M) \end{array}$$

The maps ∂_*^* count unparameterized gradient trajectories as before, so

$$\langle \partial_*^* a, b \rangle = |\check{\mathcal{M}}(a, b)|.$$

The maps $\bar{\partial}_*^*$ only count those trajectories which remain in the boundary. More precisely,

$$\langle \bar{\partial}_*^* a, b \rangle = |\check{\mathcal{M}}^\partial(a, b)|$$

where $\check{\mathcal{M}}^\partial(a, b) = \check{\mathcal{M}}(a, b) \cap (\{\gamma : \mathbb{R} \rightarrow \partial M\}/\mathbb{R})$.




FIGURE 15. Consider the surface at left embedded in \mathbb{R}^3 with the induced metric and height function. Each critical point is labeled by its type and Morse index. The directed graph at right depicts the vector spaces generated by the interior, stable, and unstable critical points together with eight linear maps between them. On the surface we have highlighted a single gradient trajectory counted by each map: four interior trajectories in black and four boundary trajectories in red. The dashed red line from s to u signifies that $\bar{\partial}_u^s$ is index-preserving. The doubled red line from u to s signifies that the $\bar{\partial}_s^u$ lowers index by 2. The other six maps lower index by 1.

The boundary, absolute, and relative Morse complexes of M are given by:

$$C(\partial M) = C^s(M) \oplus C^u(M) \quad C(M) = C^o(M) \oplus C^s(M) \quad C(M, \partial M) = C^o(M) \oplus C^u(M)$$

$$\partial_{\partial M} = \begin{bmatrix} \bar{\partial}_s^s & \bar{\partial}_s^u \\ \bar{\partial}_u^s & \bar{\partial}_u^u \end{bmatrix} \quad \partial = \begin{bmatrix} \partial_o^o & \partial_o^u \bar{\partial}_u^s \\ \partial_s^o & \bar{\partial}_s^s + \partial_s^u \bar{\partial}_u^s \end{bmatrix} \quad \partial_{M, \partial M} = \begin{bmatrix} \partial_o^o & \partial_o^u \\ \bar{\partial}_u^s \partial_s^o & \bar{\partial}_u^u + \bar{\partial}_u^s \partial_s^u \end{bmatrix}$$

The reader should compare the definition of ∂ with that of D in Figure 3. Figure 16 illustrates each complex for a particular metric and Morse function on the 2-disk.




FIGURE 16. At left, the double of the 2-disk M is symmetrically embedded in \mathbb{R}^3 . The induced metric and height function on M determine boundary, absolute, and relative Morse complexes over \mathbb{F} (so two arrows are equivalent to zero arrows). In each case, the Morse homology agrees with the singular homology of the 2-disk as expected.

For each map ∂_* , there is a map $\tilde{\partial}\delta_*$ which counts points in the boundary of the compactification of $\check{\mathcal{M}}(a, b)$. For each map $\bar{\partial}_*$, there is a map $\tilde{\partial}\delta_*$ which counts points in the boundary of the compactification of $\check{\mathcal{M}}^\partial(a, b)$. These maps should be compared with the definition of δ in Figure 3.

$$\begin{aligned} \tilde{\delta}\partial_o^o &= \partial_o^o\partial_o^o + \partial_o^u\bar{\partial}_u^s\partial_s^o & \tilde{\delta}\bar{\partial}_s^s &= \bar{\partial}_s^s\bar{\partial}_s^s + \bar{\partial}_s^u\bar{\partial}_u^s \\ \tilde{\delta}\partial_s^o &= \partial_s^o\partial_o^o + \bar{\partial}_s^s\partial_s^o + \partial_s^u\bar{\partial}_u^s\partial_s^o & \tilde{\delta}\bar{\partial}_u^s &= \bar{\partial}_u^s\bar{\partial}_s^s + \bar{\partial}_u^u\bar{\partial}_u^s \\ \tilde{\delta}\partial_o^u &= \partial_o^o\partial_o^u + \partial_o^u\bar{\partial}_u^u + \partial_o^u\bar{\partial}_u^s\partial_s^u & \tilde{\delta}\bar{\partial}_s^u &= \bar{\partial}_s^s\bar{\partial}_s^u + \bar{\partial}_s^u\bar{\partial}_u^u \\ \tilde{\delta}\partial_s^u &= \bar{\partial}_s^u + \partial_s^o\partial_o^u + \bar{\partial}_s^s\partial_s^u + \partial_s^u\bar{\partial}_u^u + \partial_s^u\bar{\partial}_u^s\partial_s^u & \tilde{\delta}\bar{\partial}_u^u &= \bar{\partial}_u^s\bar{\partial}_s^u + \bar{\partial}_u^u\bar{\partial}_u^u \end{aligned}$$

Alternatively, we may view the right-hand sides as relations, since each counts points in the boundary of a 1-dimensional space (itself a union of products of trajectory spaces). As in Figure 3, the structure equation takes the form

$$\begin{bmatrix} \tilde{\delta}\partial_o^o & (\tilde{\delta}\partial_o^u)\bar{\partial}_u^s + \partial_o^u(\tilde{\delta}\bar{\partial}_u^s) \\ \tilde{\delta}\partial_s^o & \delta\bar{\partial}_s^s + (\tilde{\delta}\partial_s^u)\bar{\partial}_u^s + \partial_s^u(\tilde{\delta}\bar{\partial}_u^s) \end{bmatrix} = \partial^2.$$

Since every term on the left vanishes, ∂ is indeed a differential. A similar argument applies to $\partial_{\partial M}$ and $\partial_{M, \partial M}$.

C. Relation to cell structures and CW-homology. On a closed manifold, a Morse-Smale function whose critical values increase with index induces a cell structure with one k -cell for each index- k critical point. The k -cell is given by the unstable manifold U_o of the critical point o_k . The Morse complex then coincides with the chain complex $C_*^{\text{cell}}(M)$ underlying CW-homology.

A similar relationship should hold for a manifold with boundary, motivated by the local picture of passing a critical level in Figure 17. We suggest that a regular function whose critical values increase with index induces a cell structure with:

- a k -cell $U_o \subset M$ for each interior critical point o_k .
- a k -cell $U_s \subset \partial M$ for each stable critical point s_k .
- a k -cell $U_u \subset M$ and a $(k-1)$ -cell $\partial U_u \subset \partial M$ for each unstable critical point u_k .

This should follow from the argument in the closed case adapted to the double \tilde{M} , though we do not attempt this here.

Assuming the above, for each unstable critical point, the $(k-1)$ -cell ∂U_u occurs with multiplicity 1 in the boundary of the k -cell U_u . There is therefore a single arrow between each pair of unstable generators in $C_*^{\text{cell}}(M)$ and a “cancellation” homotopy equivalence to a smaller complex $C'_*(M)$ generated by the interior and stable cells. The components of the differential on $C'_*(M)$ running from stable to interior generators are induced by “zig-zags” in $C_*^{\text{cell}}(M)$ involving the algebraically-cancelled pairs of unstable generators. A short exercise reveals that the complex $C'_*(M)$ is naturally identified with the Morse complex. In Figure 18, we illustrate this relationship for the example in Figure 16. A similar exercise applies to the relative and boundary complexes without cancellation. In summary:

$$\begin{aligned} C_*^{\text{cell}}(M) &\xrightarrow{\sim} C'_*(M) = C_*^{\text{morse}}(M) \\ C_*^{\text{cell}}(M, \partial M) &= C_*^{\text{morse}}(M, \partial M) \\ C_*^{\text{cell}}(\partial M) &= C_*^{\text{morse}}(\partial M) \end{aligned}$$

In this way, the rather strange Morse differentials arises naturally from the cellular perspective.




FIGURE 17. The local level-set of M just about a critical value for each type of critical point, with the mirror of M faded.




FIGURE 18. We consider the example in Figure 16 from the cellular perspective. The cellular chain complex is generated by five cells and the differential encodes the algebraic multiplicities of the attaching maps. The Morse complex results from algebraically canceling the pair of unstable cells.




FIGURE 19. The proof of the surgery exact triangle for $H\check{M}$ involves a map \check{G} defined by counting monopoles on a cobordism over a family of metrics parameterized by a pentagon. The edges of the pentagon correspond to five hypersurfaces which intersect as illustrated at top left (compare with Figure 2 in [10]). The map \check{G} corresponds to the edge labelled by the brown pentagon in the higraph, which has 7 vertices and 23 edges. In [10], the four aligned vertices are type 2 ($H\check{M}$) and the remaining three vertices are type 1 (\widehat{HM}). The path DGA on the bigraph (which need not be shown!) determines the terms of all monopole maps represented by edges in the higraph, and the structure equation implies that \check{G} satisfies the relation at the bottom of the figure (this agrees with Proposition 5.5 in [10] because the yellow and blue triangles are zero maps for geometric reasons). Note that since this higraph is tree-like, the assignment $\psi \equiv 1$ is balanced as well, leading to the surgery triangle for \widehat{HM} . The key point here is that, due to Theorem 4.11, one can justifiably work on the level of higraphs, as though all monopoles in life were irreducible after all.

REFERENCES

- [1] M. Abouzaid. A topological model for the Fukaya category of plumbings. *J. of Diff. Geom.*, 87(1):1–80, 2011.
- [2] M. Akaho. Floer’s chain complexes for Lagrangian submanifolds in symplectic manifolds with concave ends. 2003, arXiv:math/0310166v1.
- [3] J. Bloom. A link surgery spectral sequence in monopole Floer homology. *Adv. Math.*, 226(4):3216–3281, 2011.
- [4] V. Colin, P. Ghiggini, and K. Honda. Equivalence of Heegaard Floer homology and embedded contact homology via open book decompositions. *Proc. Nat. Acad. Sci.*, 108(20), 2011.
- [5] S. Devadoss. A realization of graph-associahedra. *Discrete Math.*, 309:271–276, 2009.
- [6] K. Fukaya. Floer homology for 3 manifold with boundary I. *Unpublished notes*, 1997, <http://www.math.kyoto-u.ac.jp/%7Efukaya/bdrt1.pdf>.
- [7] K. Fukaya and Y.-G. Oh. Zero-loop open strings in the cotangent bundle and Morse homotopy. *Asian J. Math*, 1(1), March 1997.
- [8] K. Fukaya, Y.-G. Oh, H. Ohta, and K. Ono. *Lagrangian intersection Floer theory - anomaly and obstruction*,. American Mathematical Society, 2010.
- [9] P. Kronheimer and T. Mrowka. *Monopoles and Three-Manifolds*. New Mathematical Monographs. Cambridge University Press, 2007.
- [10] P. Kronheimer, T. Mrowka, P. Ozsváth, and Z. Szabó. Monopoles and lens space surgeries. *Ann. of Math. (2)*, 165(2):457–546, 2007, math.GT/0310164.
- [11] C. Kutluhan, Y.-J. Lee, and C. H. Taubes. HF=HM I : Heegaard Floer homology and SeibergWitten Floer homology. 2010, arXiv:1007.1979.
- [12] Y. Lekili and T. Perutz. Fukaya categories of the torus and Dehn surgery. *Proc. Nat. Acad. Sci.*, 108(20), 2011, arXiv:1102.3160.
- [13] Y. Lekili and T. Perutz. Arithmetic mirror symmetry for the 2-torus. 2012, arXiv:1211.4632.
- [14] R. Lipshitz, P. S. Ozsváth, and D. P. Thurston. Tour of bordered Floer theory. *Proc. Nat. Acad. Sci.*, 108(20), 2011.
- [15] P. Seidel. *Fukaya categories and Picard-Lefschetz theory*. Zurich Lectures in Advanced Mathematics. European Mathematical Society, 2008.

DEPARTMENT OF MATHEMATICS, MASSACHUSETTS INSTITUTE OF TECHNOLOGY, CAMBRIDGE, USA

E-mail address: `jbloom@math.mit.edu`

URL: `http://math.mit.edu/~jbloom/`